



NVIDIA cuDNN

Developer Guide | NVIDIA Docs

Table of Contents

Chapter 1. Overview.....	1
1.1. Programming Model.....	1
1.2. GPU And Driver Requirements.....	2
1.3. Backward Compatibility And Deprecation Policy.....	2
1.4. Thread Safety.....	4
Chapter 2. Tensor Descriptor.....	5
2.1. WXYZ Tensor Descriptor.....	5
2.2. 3-D Tensor Descriptor.....	5
2.3. 4-D Tensor Descriptor.....	6
2.4. 5-D Tensor Descriptor.....	6
2.5. Fully-packed Tensors.....	6
2.6. Partially-packed Tensors.....	6
2.7. Spatially Packed Tensors.....	7
2.8. Overlapping Tensors.....	7
Chapter 3. Data Layout Formats.....	8
3.1. Example Tensor.....	8
3.2. Convolution Layouts.....	9
3.2.1. NCHW Memory Layout.....	9
3.2.2. NHWC Memory Layout.....	10
3.2.3. NC/32HW32 Memory Layout.....	12
3.3. MatMul Layouts.....	13
Chapter 4. Reproducibility (determinism).....	14
Chapter 5. Scaling Parameters.....	15
Chapter 6. Tensor Core Operations.....	17
6.1. Basics.....	17
6.2. Convolution Functions.....	18
6.2.1. Prerequisites.....	18
6.2.2. Supported Algorithms.....	18
6.2.3. Data And Filter Formats.....	18
6.3. RNN Functions.....	19
6.3.1. Prerequisites.....	19
6.3.2. Supported Algorithms.....	19
6.3.3. Data And Filter Formats.....	19
6.4. Tensor Transformations.....	20

6.4.1. FP32-to-FP16 Conversion.....	20
6.4.2. Padding.....	21
6.4.3. Folding.....	21
6.4.4. Conversion Between NCHW And NHWC.....	22
6.5. Data Transformations.....	22
6.5.1. FP16 Data.....	22
6.6. Guidelines For Good Performance On Tensor Cores.....	23
Chapter 7. Convolutions.....	24
7.1. Convolution Formulas.....	24
7.2. Grouped Convolutions.....	26
7.3. Best Practices For 3D Convolutions.....	27
7.3.1. Recommended Settings.....	28
7.3.2. Limitations.....	29
Chapter 8. Features Of RNN Functions.....	30
Chapter 9. Mixed Precision Numerical Accuracy.....	33
Chapter 10. The cuDNN Graph API.....	34
10.1. Graph API Example with Operation Fusion.....	34
10.1.1. Creating Operation and Tensor Descriptors to Specify the Graph Dataflow.....	35
10.1.2. Finalizing The Operation Graph.....	36
10.1.3. Configuring An Engine That Can Execute The Operation Graph.....	36
10.1.4. Executing The Engine.....	37
10.2. Supported Graph Patterns.....	37
10.2.1. Pre-compiled Single Operation Engines.....	37
10.2.2. Runtime Fusion Engine.....	38
10.2.2.1. Limitations.....	42
10.2.2.2. Examples of Supported Patterns.....	45
10.2.2.3. Operation specific Constraints for the Runtime Fusion Engine.....	47
10.2.3. Pre-compiled Specialized Engines.....	56
10.2.4. Mapping with Backend Descriptors.....	61
Chapter 11. Troubleshooting.....	63
11.1. Error Reporting And API Logging.....	63
11.2. FAQs.....	65
11.3. Support.....	68
Chapter 12. Acknowledgments.....	70
12.1. University of Tennessee.....	70
12.2. University of California, Berkeley.....	70
12.3. Facebook AI Research, New York.....	71

List of Figures

Figure 1. Example with N=1, C=64, H=5, W=4.....	9
Figure 2. NCHW Memory Layout.....	10
Figure 3. NHWC Memory Layout.....	11
Figure 4. NC/32HW32 Memory Layout.....	12
Figure 5. Scaling Parameters for Convolution.....	15
Figure 6. INT8 for cudnnConvolutionBiasActivationForward.....	16
Figure 7. Tensor Operation with FP32 Inputs.....	21
Figure 8. Tensor operation with FP16 inputs. The accumulation is in FP32, which could be the input for other kernel features (for example, activation/bias, beta blending, etc). The final output in this example would be FP16.....	23
Figure 9. A set of operation descriptors the user passes to the operation graph.....	35
Figure 10. The operation graph after finalization.....	36
Figure 11. ConvolutionFwd Engine.....	37
Figure 12. ConvolutionBwFilter Engine.....	38
Figure 13. ConvolutionBwData Engine.....	38
Figure 14. Graphical Representation of the following Generic Patterns Supported by the Runtime Fusion Engine.....	39
Figure 15. DAGs of cuDNN operations.....	40
Figure 16. cuDNN graph depicting DAG:Padding Mask.....	41
Figure 17. cuDNN graph depicting DAG:Softmax.....	41
Figure 18. cuDNN graph depicting DAG:Dropout.....	41
Figure 19. This example illustrates the Runtime Fusion Engine with a Single Operation.....	45
Figure 20. ConvolutionFwd Followed by a DAG with Two Operations.....	45
Figure 21. ConvolutionFwd Followed by a DAG with Three Operations.....	46
Figure 22. MatMul Preceded by a DAG with Two Operations.....	46

Figure 23. This example illustrates fusion of operations before and after the ConvolutionFwd operation. In addition we observe that the output of ConvolutionFwd can feed anywhere in g2.....	47
Figure 24. Values In the Index Tensors.....	54
Figure 25. The pre-compiled ConvBNfprop engine fuses several pointwise operations with ConvolutionFwd and GenStats.....	56
Figure 26. The ConvBNwgrad pre-compiled engine fuses several (optional) pointwise operations with ConvolutionBwFilter.....	57
Figure 27. ConvBiasAct, another pre-compiled engine, fuses ConvolutionFwd with several pointwise operations.....	57
Figure 28. The pre-compiled engine, ConvScaleBiasAct.....	58
Figure 29. The pre-compiled engine, dBNApply.....	58
Figure 30. The DualdBNApply engine.....	59
Figure 31. DgradDreluBNBwdWeight is a pre-compiled engine that can be used in conjunction with the dBNApply pattern to compute the backwards path of batch norm.....	60
Figure 32. Software stack with cuDNN.....	66

List of Tables

Table 1. Two-step, deprecation policy.....	3
Table 2. Convolution terms.....	24
Table 3. Recommended settings while performing 3D convolutions for cuDNN.....	28
Table 4. Limitations per Generic Pattern.....	42
Table 5. Limitations per Generic Pattern.....	43
Table 6. Layout Requirements per Pattern.....	44
Table 7. Tensor Attributes for all Three Operations.....	48
Table 8. Constraints for all Three Operations.....	48
Table 9. I/O Tensors Alignment Requirements.....	48
Table 10. Batch Size Requirements Per Operation.....	49
Table 11. Recommended compute type for FP8 tensor computations for Hopper architecture.....	49
Table 12. Constraints for MatMul Operations.....	50
Table 13. MatMul Alignment Requirements.....	50
Table 14. Constraints for Pointwise Operations.....	50
Table 15. Constraints for GenStats Operations.....	51
Table 16. Constraints for Reduction Operations.....	52
Table 17. Supported Reduction Patterns.....	52
Table 18. Specific Restrictions for the Downsampling Modes.....	53
Table 19. Specific Restrictions for Upsampling Mode CUDNN_RESAMPLE_BILINEAR.....	53
Table 20. Specific Restrictions for the Backwards Downsampling Modes.....	55
Table 21. Notations and Backend Descriptors.....	61
Table 22. API Logging Using Environment Variables.....	65

Chapter 1. Overview

NVIDIA® CUDA® Deep Neural Network Library (cuDNN) is a GPU-accelerated library of primitives for deep neural networks. It provides highly tuned implementations of routines arising frequently in DNN applications:

- ▶ Convolution forward and backward, including cross-correlation
- ▶ Matrix multiplication
- ▶ Pooling forward and backward
- ▶ Softmax forward and backward
- ▶ Neuron activations forward and backward: `relu`, `tanh`, `sigmoid`, `elu`, `gelu`, `softplus`, `swish`
- ▶ Arithmetic, mathematical, relational and logical pointwise operations
- ▶ Tensor transformation functions
- ▶ LRN, LCN and batch normalization forward and backward

cuDNN convolution routines aim for a performance that is competitive with the fastest GEMM (matrix multiply)-based implementations of such routines while using significantly less memory.

cuDNN features include customizable data layouts, supporting flexible dimension ordering, striding, and subregions for the 4D tensors used as inputs and outputs to all of its routines. This flexibility allows easy integration into any neural network implementation and avoids the input/output transposition steps sometimes necessary with GEMM-based convolutions.

cuDNN offers a context-based API that allows for easy multithreading and (optional) interoperability with CUDA® streams.

1.1. Programming Model

The cuDNN library exposes a host API but assumes that for operations using the GPU, the necessary data is directly accessible from the device.

An application using cuDNN must initialize a handle to the library context by calling [`cudaDnnCreate\(\)`](#). This handle is explicitly passed to every subsequent library function that operates on GPU data. Once the application finishes using cuDNN, it can release the resources associated with the library handle using [`cudaDnnDestroy\(\)`](#). This approach allows

the user to explicitly control the library's functioning when using multiple host threads, GPUs and CUDA streams.

For example, an application can use [cudaSetDevice](#) (prior to creating a cuDNN handle) to associate different devices with different host threads, and in each of those host threads, create a unique cuDNN handle that directs the subsequent library calls to the device associated with it. Therefore, the cuDNN library calls made with different handles will automatically run on different devices.

The device associated with a particular cuDNN context is assumed to remain unchanged between the corresponding `cudaCreate()` and `cudaDestroy()` calls. In order for the cuDNN library to use a different device within the same host thread, the application must set the new device to be used by calling `cudaSetDevice()` and then create another cuDNN context, which will be associated with the new device, by calling `cudaCreate()`.

cuDNN API Compatibility

Beginning in cuDNN 7, the binary compatibility of a patch and minor releases is maintained as follows:

- ▶ Any patch release x.y.z is forward or backward-compatible with applications built against another cuDNN patch release x.y.w (meaning, of the same major and minor version number, but having w!=z).
- ▶ cuDNN minor releases beginning with cuDNN 7 are binary backward-compatible with applications built against the same or earlier patch release (meaning, an application built against cuDNN 7.x is binary compatible with cuDNN library 7.y, where y>=x).
- ▶ Applications compiled with a cuDNN version 7.y are not guaranteed to work with 7.x release when y > x.

1.2. GPU And Driver Requirements

For the latest compatibility software versions of the OS, CUDA, the CUDA driver, and the NVIDIA hardware, see the [NVIDIA cuDNN Support Matrix](#).

1.3. Backward Compatibility And Deprecation Policy

cuDNN version 8 introduces a new API deprecation policy to enable a faster pace of innovation.

The old deprecation policy required three major library releases to complete an API update. During this process, the original function name was first assigned to the legacy API, and then to the revised API, depending on the library version. The user wishing to migrate to the new API version had to update his or her code twice. In the first update, the original call `fnn()` had to be changed to `fnn_vN()`, where N is the new major cuDNN version. After the next major cuDNN release, the `fnn_vN()` function had to be renamed

back as `foo()`. Clearly, the above process could be difficult for code maintenance, especially when many functions are upgraded.

A streamlined, two-step, deprecation policy will be used for all API changes starting with cuDNN version 8. Let us explain the process using two subsequent, major cuDNN releases, version 8 and 9:

Table 1. Two-step, deprecation policy

cuDNN version	Explanation
Major release 8	The updated API is introduced as <code>foo_v8()</code> . The deprecated API <code>foo()</code> is kept unchanged to maintain backward compatibility until the next major release.
Major release 9	The deprecated API <code>foo()</code> is permanently removed and its name is not reused. The <code>foo_v8()</code> function supersedes the retired call <code>foo()</code> .

If the existing API needs to be updated, a new function flavor is introduced with the `_v` tag followed by the current, major cuDNN version. In the next major release, the deprecated function is removed, and its name is never reused. A brand-new API is first introduced without the `_v` tag.

The revised depreciation scheme allows us to retire the legacy API in just one major release. Similarly to the previous API deprecation policy, the user is able to compile the legacy code without any changes using the next major release of the cuDNN library. The backward compatibility ends when another major cuDNN release is introduced.

The updated function name embeds the information in which the cuDNN version of the API call was modified. As a result, the API changes will be easier to track and document.

The new deprecation policy is applied also to pending API changes from previous cuDNN releases. For example, according to the old deprecation policy, `cudaSetRNNDescriptor_v6()` should be removed in cuDNN version 8 and the upgraded call `cudaSetRNNDescriptor()` with the same arguments and behavior should be kept. Instead, the new deprecation policy is applied to this case and the tagged function is kept.

Prototypes of deprecated functions will be prepended in cuDNN version 8 headers using the `CUDNN_DEPRECATED` macro. When the `-DCUDNN_WARN_DEPRECATED` switch is passed to the compiler, any deprecated function call in the user's code will emit a compiler warning, for example:

```
warning: 'cudaStatus_t cudaSetRNNMatrixMathType(cudaRNNDescriptor_t, cudaMathType_t)' is deprecated [-Wdeprecated-declarations]
```

Or

```
warning C4996: 'cudaSetRNNMatrixMathType': was declared deprecated
```

The above warnings are disabled by default to avoid potential build breaks in software setups where compiler warnings are treated as errors.

Note that the simple swapping of older cuDNN version 7 shared library files will not work with the cuDNN version 8 release. The user source code needs to be recompiled from scratch with the cuDNN version 8 headers and linked with the version 8 libraries.

1.4. Thread Safety

The cuDNN library is thread-safe. Its functions can be called from multiple host threads, so long as the threads do not share the same cuDNN handle simultaneously.

When creating a per-thread cuDNN handle, it is recommended that a single synchronous call of [cudnnCreate\(\)](#) be made first before each thread creates its own handle asynchronously.

Per [cudnnCreate\(\)](#), for multi-threaded applications that use the same device from different threads, the recommended programming model is to create one (or a few, as is convenient) cuDNN handles per thread and use that cuDNN handle for the entire life of the thread.

Chapter 2. Tensor Descriptor

The cuDNN library describes data holding images, videos and any other data with contents with a generic n-D tensor defined with the following parameters:

- ▶ a dimension `nbDims` from 3 to 8
- ▶ a data type (32-bit floating-point, 64 bit-floating point, 16-bit floating-point...)
- ▶ `dimA` integer array defining the size of each dimension
- ▶ `strideA` integer array defining the stride of each dimension (for example, the number of elements to add to reach the next element from the same dimension)

The first dimension of the tensor defines the batch size `n`, and the second dimension defines the number of features maps `c`. This tensor definition allows, for example, to have some dimensions overlapping each other within the same tensor by having the stride of one dimension smaller than the product of the dimension and the stride of the next dimension. In cuDNN, unless specified otherwise, all routines will support tensors with overlapping dimensions for forward-pass input tensors, however, dimensions of the output tensors cannot overlap. Even though this tensor format supports negative strides (which can be useful for data mirroring), cuDNN routines do not support tensors with negative strides unless specified otherwise.

2.1. WXYZ Tensor Descriptor

Tensor descriptor formats are identified using acronyms, with each letter referencing a corresponding dimension. In this document, the usage of this terminology implies:

- ▶ all the strides are strictly positive
- ▶ the dimensions referenced by the letters are sorted in decreasing order of their respective strides

2.2. 3-D Tensor Descriptor

A 3-D tensor is commonly used for matrix multiplications, with three letters: B, M, and N. B represents the batch size (for batch GEMM, set to 1 for single GEMM), M represents the number of rows, and N represents the number of columns. Refer to the [MatMul](#) operation for more information.

2.3. 4-D Tensor Descriptor

A 4-D tensor descriptor is used to define the format for batches of 2D images with 4 letters: N, C, H, W for respectively the batch size, the number of feature maps, the height and the width. The letters are sorted in decreasing order of the strides. The commonly used 4-D tensor formats are:

- ▶ NCHW
- ▶ NHWC
- ▶ CHWN

2.4. 5-D Tensor Descriptor

A 5-D tensor descriptor is used to define the format of the batch of 3D images with 5 letters: N, C, D, H, W for respectively the batch size, the number of feature maps, the depth, the height, and the width. The letters are sorted in decreasing order of the strides. The commonly used 5-D tensor formats are called:

- ▶ NCDHW
- ▶ NDHWC
- ▶ CDHWN

2.5. Fully-packed Tensors

A tensor is defined as *XYZ-fully-packed* if and only if:

- ▶ the number of tensor dimensions is equal to the number of letters preceding the *fully-packed* suffix.
- ▶ the stride of the i -th dimension is equal to the product of the $(i+1)$ -th dimension by the $(i+1)$ -th stride.
- ▶ the stride of the last dimension is 1.

2.6. Partially-packed Tensors

The partially *XYZ-packed* terminology only applies in the context of a tensor format described with a superset of the letters used to define a partially-packed tensor. A $WXYZ$ tensor is defined as *XYZ-packed* if and only if:

- ▶ The strides of all dimensions NOT referenced in the *-packed* suffix are greater or equal to the product of the next dimension by the next stride.
- ▶ The stride of each dimension referenced in the *-packed* suffix in position i is equal to the product of the $(i+1)$ -st dimension by the $(i+1)$ -st stride.

- If the last tensor's dimension is present in the `-packed` suffix, its stride is 1.

For example, an `NHWC` tensor `WC-packed` means that the `c_stride` is equal to 1 and `w_stride` is equal to `c_dim x c_stride`. In practice, the `-packed` suffix is usually applied to the minor dimensions of a tensor but can be applied to only the major dimensions; for example, an `NCHW` tensor that is only `N-packed`.

2.7. Spatially Packed Tensors

Spatially-packed tensors are defined as partially-packed in spatial dimensions. For example, a spatially-packed 4D tensor would mean that the tensor is either `NCHW HW-packed` or `CNHW HW-packed`.

2.8. Overlapping Tensors

A tensor is defined to be overlapping if iterating over a full range of dimensions produces the same address more than once. In practice an overlapped tensor will have `stride[i-1] < stride[i]*dim[i]` for some of the `i` from `[1,nbDims]` interval.

Chapter 3. Data Layout Formats

This section describes how cuDNN tensors are arranged in memory according to several data layout formats.

The recommended way to specify the layout format of a tensor is by setting its strides accordingly. For compatibility with the v7 API, a subset of the layout formats can also be configured through the [cudnnTensorFormat_t](#) enum in the *cuDNN API Reference*. The enum is only supplied for legacy reasons and is deprecated.

3.1. Example Tensor

Consider a batch of images with the following dimensions:

- ▶ N is the batch size; 1.
- ▶ C is the number of feature maps (i.e., number of channels); 64.
- ▶ H is the image height; 5.
- ▶ W is the image width; 4.

To keep the example simple, the image pixel elements are expressed as a sequence of integers, 0, 1, 2, 3, and so on. See [Figure 1](#).

Figure 1. Example with N=1, C=64, H=5, W=4.

EXAMPLE

N = 1

C = 64

H = 5

W = 4

c = 0

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15
16	17	18	19

c = 1

20	21	22	23
24	25	26	27
28	29	30	31
32	33	34	35
36	37	38	39

...

c = 2

40	41	42	43
44	45	46	47
48	49	50	51
52	53	54	55
56	57	58	59

c = 30

600	601	602	603
604	605	606	607
608	609	610	611
612	613	614	615
616	617	618	619

c = 31

620	621	622	623
624	625	626	627
628	629	630	631
632	633	634	635
636	637	638	639

...

c = 32

640	641	642	643
644	645	646	647
648	649	650	651
652	653	654	655
656	657	658	659

c = 62

1240	1241	1242	1243
1244	1245	1246	1247
1248	1249	1250	1251
1252	1253	1254	1255
1256	1257	1258	1259

c = 63

1260	1261	1262	1263
1264	1265	1266	1267
1268	1269	1270	1271
1272	1273	1274	1275
1276	1277	1278	1279

In the following subsections, we'll use the above example to demonstrate the different layout formats.

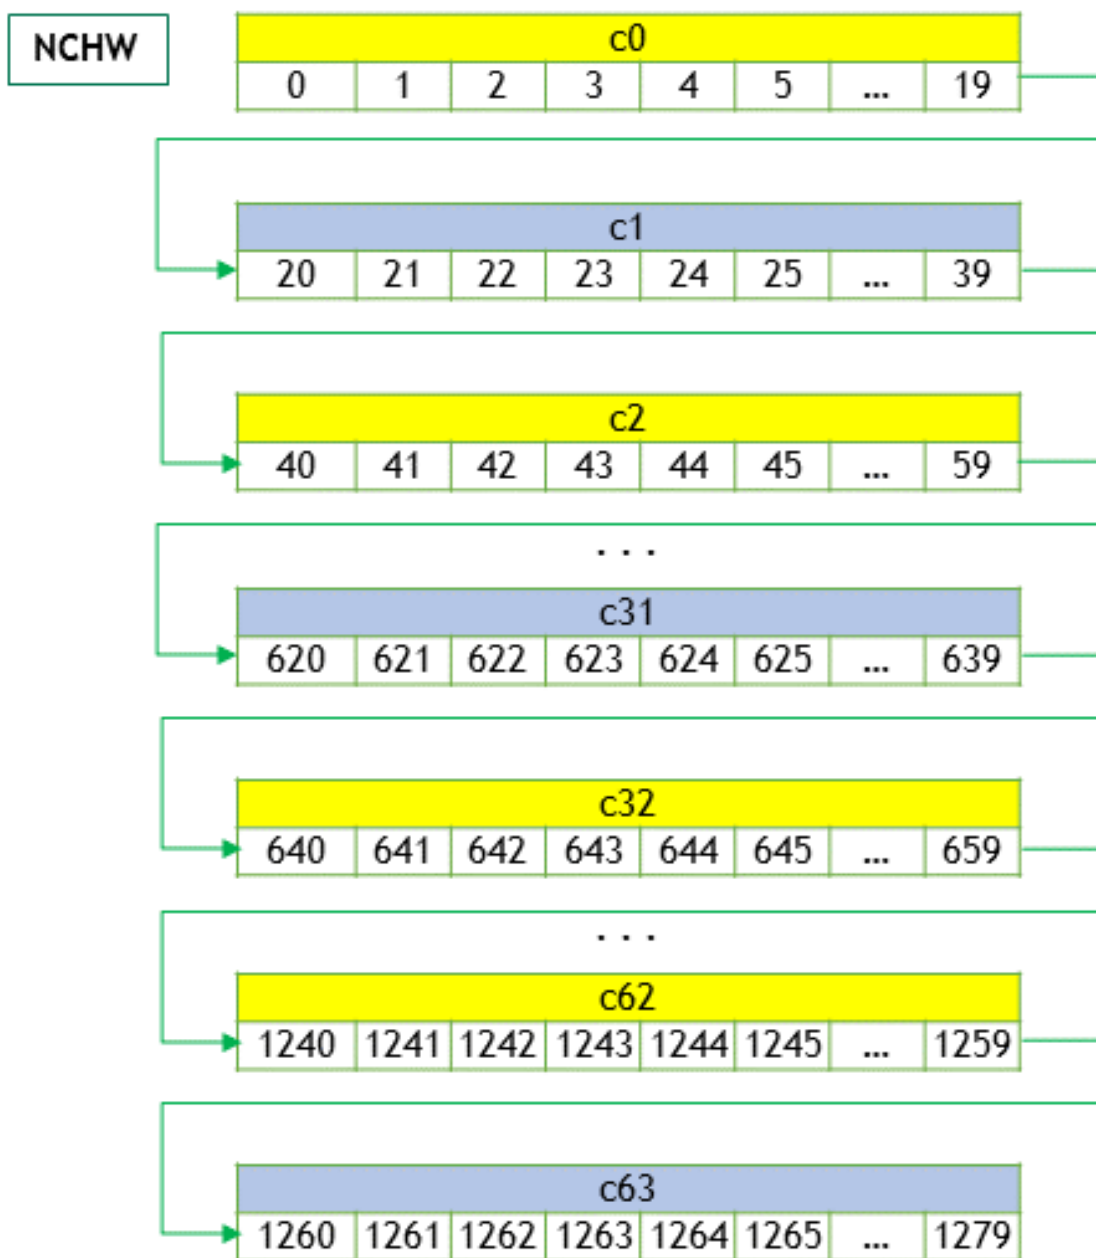
3.2. Convolution Layouts

3.2.1. NCHW Memory Layout

The above 4D tensor is laid out in the memory in the NCHW format as below:

1. Beginning with the first channel (c=0), the elements are arranged contiguously in row-major order.
2. Continue with second and subsequent channels until the elements of all the channels are laid out. Refer to [Figure 2](#).
3. Proceed to the next batch (if N is > 1).

Figure 2. NCHW Memory Layout



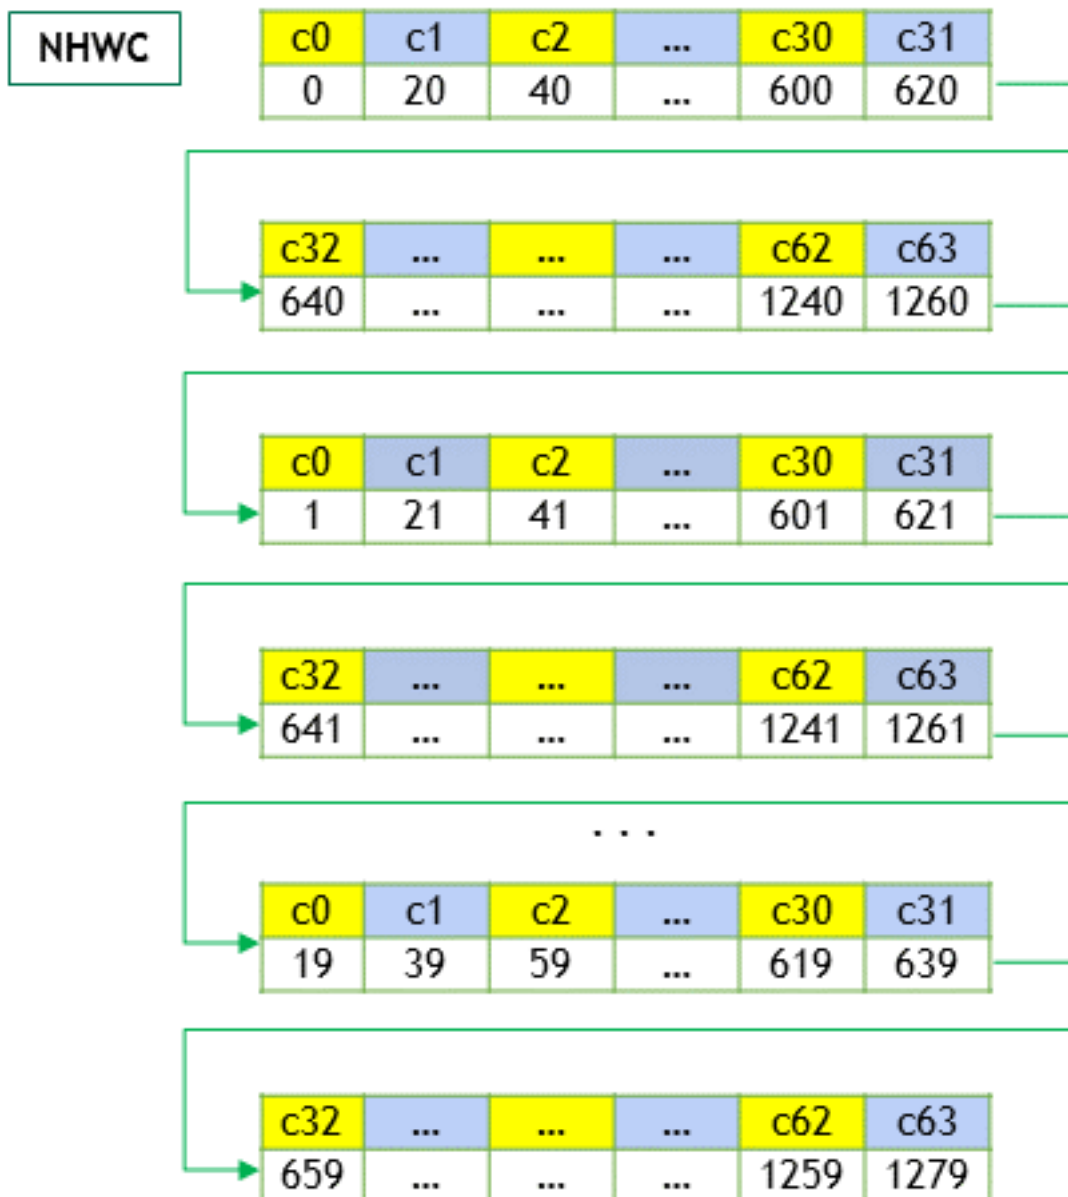
3.2.2. NHWC Memory Layout

For the NHWC memory layout, the corresponding elements in all the C channels are laid out first, as below:

1. Begin with the first element of channel 0, then proceed to the first element of channel 1, and so on, until the first elements of all the C channels are laid out.

2. Next, select the second element of channel 0, then proceed to the second element of channel 1, and so on, until the second element of all the channels are laid out.
3. Follow the row-major order of channel 0 and complete all the elements. Refer to [Figure 3](#).
4. Proceed to the next batch (if N is > 1).

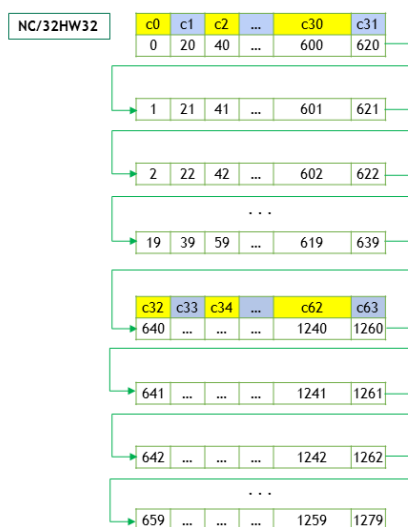
Figure 3. NHWC Memory Layout



3.2.3. NC/32HW32 Memory Layout

The NC/32HW32 is similar to NHWC, with a key difference. For the NC/32HW32 memory layout, the 64 channels are grouped into two groups of 32 channels each - first group consisting of channels c_0 through c_{31} , and the second group consisting of channels c_{32} through c_{63} . Then each group is laid out using the NHWC format. See [Figure 4](#).

Figure 4. NC/32HW32 Memory Layout



For the generalized NC/xHWx layout format, the following observations apply:

- ▶ Only the channel dimension, c , is grouped into x channels each.
- ▶ When $x = 1$, each group has only one channel. Hence, the elements of one channel (i.e., one group) are arranged contiguously (in the row-major order), before proceeding to the next group (i.e., next channel). This is the same as NCHW format.
- ▶ When $x = c$, then NC/xHWx is identical to NHWC, i.e., the entire channel depth c is considered as a single group. The case $x = c$ can be thought of as vectorizing the entire c dimension as one big vector, laying out all the c s, followed by the remaining dimensions, just like NHWC.
- ▶ The tensor format `CUDNN_TENSOR_NCHW_VECT_C` can also be interpreted in the following way: The NCHW INT8x32 format is really $N \times (C/32) \times H \times W \times 32$ (32 c s for every w), just as the NCHW INT8x4 format is $N \times (C/4) \times H \times W \times 4$ (4 c s for every w). Hence the `VECT_C` name - each w is a vector (4 or 32) of c s.

3.3. MatMul Layouts

As discussed in [3-D Tensor Descriptor](#), MatMul uses 3D tensors, described using BMN dimensions. The layout can be specified through the following strides. The following are two examples of recommended layouts:

- ▶ Packed Row-major: dim [B,M,N] with stride [MN, N, 1], or
- ▶ Packed Column-major: dim [B,M,N] with stride [MN, 1, M]

Unpacked layouts for 3-D tensors are supported as well, but their support surface is more ragged.

Chapter 4. Reproducibility (determinism)

By design, most of cuDNN's routines from a given version generate the same bit-wise results across runs when executed on GPUs with the same architecture. There are some exceptions. For example, the following routines do not guarantee reproducibility across runs, even on the same architecture, because they use atomic operations in a way that introduces truly random floating point rounding errors:

- ▶ `cudaConvolutionBackwardFilter` when `CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0` or `CUDNN_CONVOLUTION_BWD_FILTER_ALGO_3` is used
- ▶ `cudaConvolutionBackwardData` when `CUDNN_CONVOLUTION_BWD_DATA_ALGO_0` is used
- ▶ `cudaPoolingBackward` when `CUDNN_POOLING_MAX` is used
- ▶ `cudaSpatialTfSamplerBackward`
- ▶ `cudaCTCLoss` and `cudaCTCLoss_v8` when `CUDNN CTC_LOSS_ALGO_NON_DETERMINISTIC` is used

Across different architectures, no cuDNN routines guarantee bit-wise reproducibility. For example, there is no guarantee of bit-wise reproducibility when comparing the same routine run on NVIDIA Volta™ and NVIDIA Turing™, or NVIDIA Turing and NVIDIA Ampere Architecture.

Chapter 5. Scaling Parameters

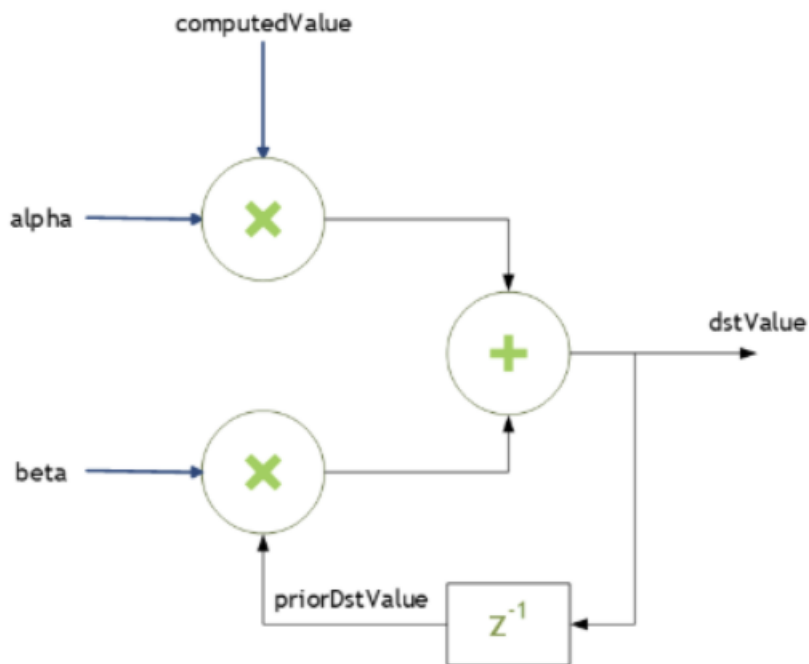
Many cuDNN routines like [cudnnConvolutionForward\(\)](#) accept pointers in host memory to scaling factors `alpha` and `beta`. These scaling factors are used to blend the computed values with the prior values in the destination tensor as follows (refer to [Figure 5](#)):

```
dstValue = alpha*computedValue + beta*priorDstValue
```



Note: The `dstValue` is written to after being read.

Figure 5. Scaling Parameters for Convolution



When `beta` is zero, the output is not read and may contain uninitialized data (including NaN).

These parameters are passed using a host memory pointer. The storage data types for `alpha` and `beta` are:

- `float` for HALF and FLOAT tensors, and
- `double` for DOUBLE tensors.



Note: For improved performance use `beta = 0.0`. Use a non-zero value for `beta` only when you need to blend the current output tensor values with the prior values of the output tensor.

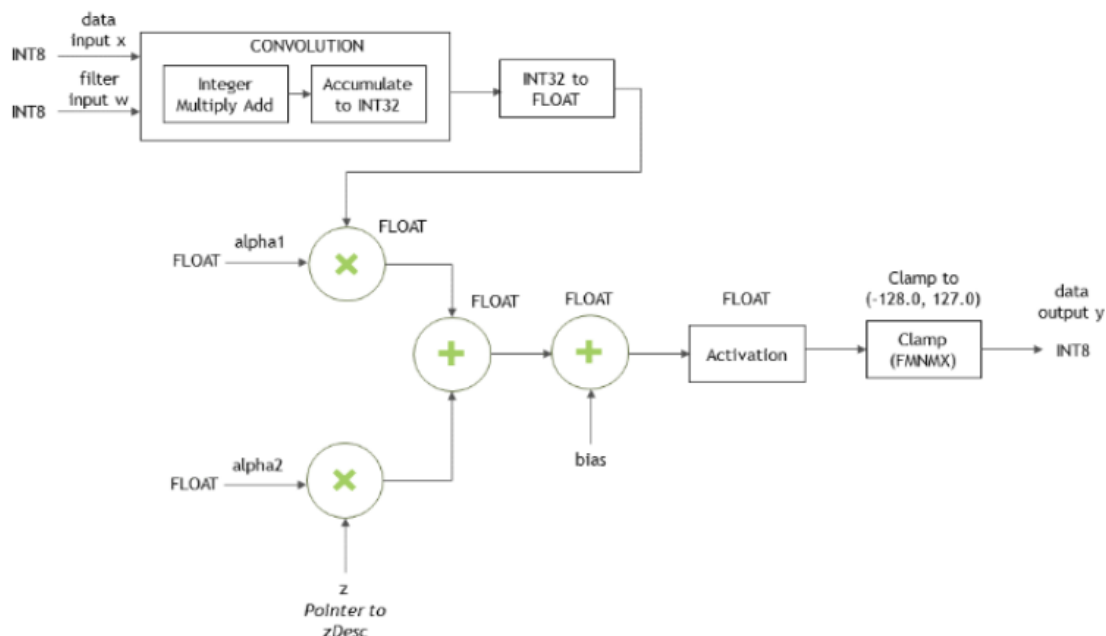
Type Conversion

When the data input x , the filter input w and the output y are all in INT8 data type, the function `cudnnConvolutionBiasActivationForward()` will perform the type conversion as shown in [Figure 6](#):



Note: Accumulators are 32-bit integers that wrap on overflow.

Figure 6. INT8 for `cudnnConvolutionBiasActivationForward`



Chapter 6. Tensor Core Operations

The cuDNN v7 library introduced the acceleration of compute-intensive routines using Tensor Core hardware on supported GPU SM versions. Tensor Core operations are supported beginning with the NVIDIA Volta GPU.

6.1. Basics

Tensor Core operations accelerate matrix math operations; cuDNN uses Tensor Core operations that accumulate into FP16, FP32, and INT32 values. Setting the math mode to `CUDNN_TENSOR_OP_MATH` via the [`cudaMathType_t`](#) enumerator indicates that the library will use Tensor Core operations. This enumerator specifies the available options to enable the Tensor Core and should be applied on a per-routine basis.

The default math mode is `CUDNN_DEFAULT_MATH`, which indicates that the Tensor Core operations will be avoided by the library. Because the `CUDNN_TENSOR_OP_MATH` mode uses the Tensor Cores, it is possible that these two modes generate slightly different numerical results due to different sequencing of the floating-point operations.

For example, the result of multiplying two matrices using Tensor Core operations is very close, but not always identical, to the result achieved using a sequence of scalar floating-point operations. For this reason, the cuDNN library requires an explicit user opt-in before enabling the use of Tensor Core operations.

However, experiments with training common deep learning models show negligible differences between using Tensor Core operations and scalar floating point paths, as measured by both the final network accuracy and the iteration count to convergence. Consequently, the cuDNN library treats both modes of operation as functionally indistinguishable and allows for the scalar paths to serve as legitimate fallbacks for cases in which the use of Tensor Core operations is unsuitable.

Kernels using Tensor Core operations are available for:

- ▶ Convolutions
- ▶ RNNs
- ▶ Multi-Head Attention

See also [Training with Mixed Precision](#).

6.2. Convolution Functions

6.2.1. Prerequisites

For the supported GPUs, the Tensor Core operations will be triggered for convolution functions only when [`cudaSetConvolutionMathType\(\)`](#) is called on the appropriate convolution descriptor by setting the `mathType` to `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION`.

6.2.2. Supported Algorithms

When the prerequisite is met, the below convolution functions can be run as Tensor Core operations:

- ▶ [`cudaConvolutionForward\(\)`](#)
- ▶ [`cudaConvolutionBackwardData\(\)`](#)
- ▶ [`cudaConvolutionBackwardFilter\(\)`](#)

See the table below for supported algorithms:

Supported Convolution Function	Supported Algos
<code>cudaConvolutionForward</code>	<code>CUDNN_CONVOLUTION_FWD_ALGO_IMPLICIT_PRECOMP_GEMM</code> <code>CUDNN_CONVOLUTION_FWD_ALGO_WINOGRAD_NONFUSED</code>
<code>cudaConvolutionBackwardData</code>	<code>CUDNN_CONVOLUTION_BWD_DATA_ALGO_1</code> <code>CUDNN_CONVOLUTION_BWD_DATA_ALGO_WINOGRAD_NONFUSED</code>
<code>cudaConvolutionBackwardFilter</code>	<code>CUDNN_CONVOLUTION_BWD_FILTER_ALGO_1</code> <code>CUDNN_CONVOLUTION_BWD_FILTER_ALGO_WINOGRAD_NONFUSED</code>

6.2.3. Data And Filter Formats

The cuDNN library may use padding, folding, and NCHW-to-NHWC transformations to call the Tensor Core operations. See [Tensor Transformations](#).

For algorithms other than `*_ALGO_WINOGRAD_NONFUSED`, when the following requirements are met, the cuDNN library will trigger the Tensor Core operations:

- ▶ Input, filter, and output descriptors (`xDesc`, `yDesc`, `wDesc`, `dxDesc`, `dyDesc` and `dwDesc` as applicable) are of the `dataType = CUDNN_DATA_HALF` (i.e., FP16). For FP32 `dataType` see [FP32-to-FP16 Conversion](#).
- ▶ The number of input and output feature maps (i.e., channel dimension `c`) is a multiple of 8. When the channel dimension is not a multiple of 8, see [Padding](#).
- ▶ The filter is of type `CUDNN_TENSOR_NCHW` or `CUDNN_TENSOR_NHWC`.

- ▶ If using a filter of type `CUDNN_TENSOR_NHWC`, then the input, filter, and output data pointers (`X`, `Y`, `W`, `dX`, `dY`, and `dW` as applicable) are aligned to 128-bit boundaries.

6.3. RNN Functions

6.3.1. Prerequisites

Tensor Core operations are triggered for these RNN functions only when `cudaSetRNNMatrixMathType()` is called on the appropriate RNN descriptor setting `mathType` to `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION`.

6.3.2. Supported Algorithms

When the above prerequisite is met, the RNN functions below can be run as Tensor Core operations:

- ▶ `cudaRnnForwardInference()`
- ▶ `cudaRnnForwardTraining()`
- ▶ `cudaRnnBackwardData()`
- ▶ `cudaRnnBackwardWeights()`
- ▶ `cudaRnnForwardInferenceEx()`
- ▶ `cudaRnnForwardTrainingEx()`
- ▶ `cudaRnnBackwardDataEx()`
- ▶ `cudaRnnBackwardWeightsEx()`
- ▶ `cudaRnnForward()`
- ▶ `cudaRnnBackwardData_v8()`
- ▶ `cudaRnnBackwardWeights_v8()`

See the table below for the supported algorithms:

RNN Function	Support Algos
All RNN functions that support Tensor Core operations.	<code>CUDNN_RNN_ALGO_STANDARD</code> <code>CUDNN_RNN_ALGO_PERSIST_STATIC</code>

6.3.3. Data And Filter Formats

When the following requirements are met, then the cuDNN library triggers the Tensor Core operations:

- ▶ For `algo = CUDNN_RNN_ALGO_STANDARD`:
 - ▶ The hidden state size, input size, and the batch size is a multiple of 8.

- ▶ All user-provided tensors, workspace, and reserve space are aligned to 128-bit boundaries.
- ▶ For FP16 input/output, the `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` is selected.
- ▶ For FP32 input/output, `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` is selected.
- ▶ For `algo = CUDNN_RNN_ALGO_PERSIST_STATIC`:
 - ▶ The hidden state size and the input size is a multiple of 32.
 - ▶ The batch size is a multiple of 8.
 - ▶ If the batch size exceeds 96 (for forward training or inference) or 32 (for backward data), then the batch size constraints may be stricter, and large power-of-two batch sizes may be needed.
 - ▶ All user-provided tensors, workspace, and reserve space are aligned to 128-bit boundaries.
 - ▶ For FP16 input/output, `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` is selected.
 - ▶ For FP32 input/output, `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` is selected.

See also [Features Of RNN Functions](#).


6.4. Tensor Transformations

A few functions in the cuDNN library will perform transformations such as folding, padding, and NCHW-to-NHWC conversion while performing the actual function operation. See below.

6.4.1. FP32-to-FP16 Conversion

The cuDNN API allows the user to specify that FP32 input data may be copied and converted to FP16 data internally to use Tensor Core operations for potentially improved performance. This can be achieved by selecting `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` enum for `cudnnMathType_t`. In this mode, the FP32 tensors are internally down-converted to FP16, the Tensor Op math is performed, and finally up-converted to FP32 as outputs. See [Figure 7](#).

Figure 7. Tensor Operation with FP32 Inputs



For Convolutions

For convolutions, the FP32-to-FP16 conversion can be achieved by passing the `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` enum value to the `cudaSetConvolutionMathType()` call.

```
// Set the math type to allow cuDNN to use Tensor Cores:
checkCudnnErr(cudaSetConvolutionMathType(cudaConvDesc,
    CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION));
```

For RNNs

For RNNs, the FP32-to-FP16 conversion can be achieved by passing the `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` enum value to the `cudaSetRNNMatrixMathType()` call to allow FP32 data to be converted for use in RNNs.

```
// Set the math type to allow cuDNN to use Tensor Cores:
checkCudnnErr(cudaSetRNNMatrixMathType(cudaRnnDesc,
    CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION));
```

6.4.2. Padding

For packed NCHW data, when the channel dimension is not a multiple of 8, then the cuDNN library will pad the tensors as needed to enable Tensor Core operations. This padding is automatic for packed NCHW data in both the `CUDNN_TENSOR_OP_MATH` and the `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` cases.

6.4.3. Folding

In the folding operation, the cuDNN library implicitly performs the formatting of input tensors and saves the input tensors in an internal workspace. This can lead to an acceleration of the call to Tensor Cores.

With folding or channel-folding, cuDNN can implicitly format the input tensors within an internal workspace to accelerate the overall calculation. Performing this transformation for the user often allows cuDNN to use kernels with restrictions on convolution stride to support a strided convolution problem.

6.4.4. Conversion Between NCHW And NHWC

Tensor Cores require that the tensors be in the NHWC data layout. Conversion between NCHW and NHWC is performed when the user requests Tensor Op math. However, as stated in [Basics](#), a request to use Tensor Cores is just that, a request and Tensor Cores may not be used in some cases. The cuDNN library converts between NCHW and NHWC if and only if Tensor Cores are requested and are actually used.

If your input (and output) are NCHW, then expect a layout change.

Non-Tensor Op convolutions will not perform conversions between NCHW and NHWC.

In very rare and difficult-to-qualify cases that are a complex function of padding and filter sizes, it is possible that Tensor Ops is not enabled. In such cases, users can pre-pad to enable the Tensor Ops path.

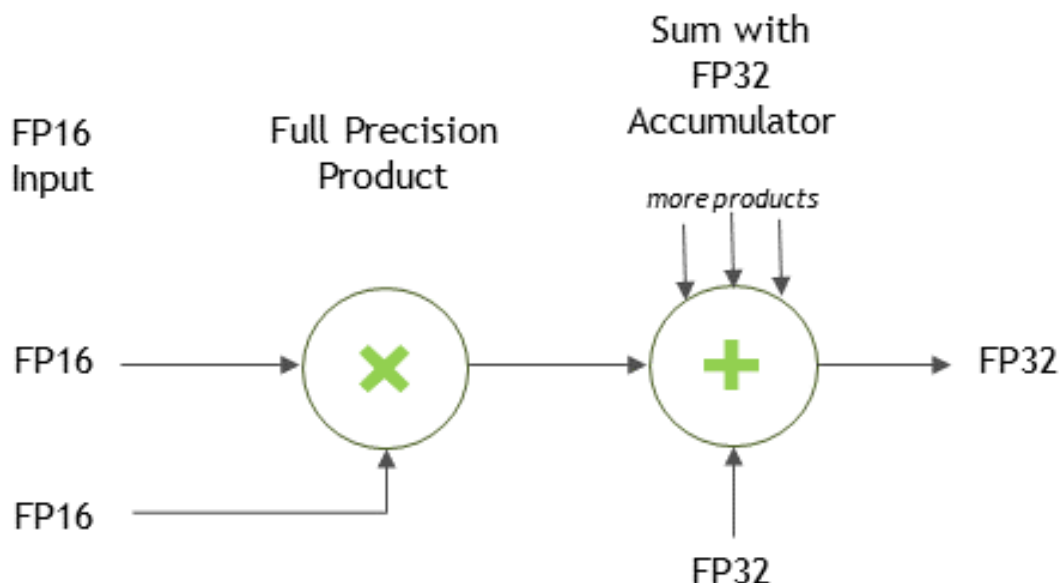
6.5. Data Transformations

6.5.1. FP16 Data

For FP16 data, Tensor Cores operate on FP16 input, output in FP16, and may accumulate in FP16 or FP32. The FP16 multiply leads to a full-precision result that is accumulated in FP32 operations with the other products in a given dot product for a matrix with $m \times n \times k$ dimensions. See [Figure 8](#).

For an FP32 accumulation, with FP16 output, the output of the accumulator is down-converted to FP16. Generally, the accumulation type is of greater or equal precision to the output type.

Figure 8. Tensor operation with FP16 inputs. The accumulation is in FP32, which could be the input for other kernel features (for example, activation/bias, beta blending, etc). The final output in this example would be FP16.



6.6. Guidelines For Good Performance On Tensor Cores

For a deep learning compiler, the following are the key guidelines:

- ▶ Make sure that the convolution operation is eligible for Tensor Cores by avoiding any combinations of large padding and large filters.
- ▶ Transform the inputs and filters to NHWC, pre-pad channel and batch size to be a multiple of 8.
- ▶ Make sure that all user-provided tensors, workspace, and reserve space are aligned to 128-bit boundaries. Note that 1024-bit alignment may deliver better performance.

Chapter 7. Convolutions

The convolution functions are:

- ▶ [cudnnConvolutionBackwardData\(\)](#)
- ▶ [cudnnConvolutionBiasActivationForward\(\)](#)
- ▶ [cudnnConvolutionForward\(\)](#)
- ▶ [cudnnConvolutionBackwardBias\(\)](#)
- ▶ [cudnnConvolutionBackwardFilter\(\)](#)

7.1. Convolution Formulas

This section describes the various convolution formulas implemented in `convolution` functions for the `cudnnConvolutionForward()` path.

The convolution terms described in the table below apply to all the convolution formulas that follow.

Table 2. Convolution terms

Term	Description
x	Input (image) Tensor
w	Weight Tensor
y	Output Tensor
n	Current Batch Size
c	Current Input Channel
C	Total Input Channels
H	Input Image Height
W	Input Image Width
k	Current Output Channel
K	Total Output Channels
p	Current Output Height Position
q	Current Output Width Position

Term	Description
G	Group Count
pad	Padding Value
u	Vertical Subsample Stride (along Height)
v	Horizontal Subsample Stride (along Width)
dil_h	Vertical Dilation (along Height)
dil_w	Horizontal Dilation (along Width)
r	Current Filter Height
R	Total Filter Height
s	Current Filter Width
S	Total Filter Width
C_g	$\frac{C}{G}$
K_g	$\frac{K}{G}$

Convolution (convolution mode set to CUDNN_CROSS_CORRELATION)

$$y_{n, k, p, q} = \sum_c^C \sum_r^R \sum_s^S x_{n, c, p+r, q+s} \times w_{k, c, r, s}$$

Convolution with Padding

$$x_{<0, <0} = 0$$

$$x_{>H, >W} = 0$$

$$y_{n, k, p, q} = \sum_c^C \sum_r^R \sum_s^S x_{n, c, p+r-pad, q+s-pad} \times w_{k, c, r, s}$$

Convolution with Subsample-Striding

$$y_{n, k, p, q} = \sum_c^C \sum_r^R \sum_s^S x_{n, c, (p*u) + r, (q*v) + s} \times w_{k, c, r, s}$$

Convolution with Dilation

$$y_{n, k, p, q} = \sum_c^C \sum_r^R \sum_s^S x_{n, c, p + (r*dil_h), q + (s*dil_w)} \times w_{k, c, r, s}$$

Convolution (convolution mode set to CUDNN_CONVOLUTION)

$$y_{n, k, p, q} = \sum_c^C \sum_r^R \sum_s^S x_{n, c, p + r, q + s} \times w_{k, c, R-r-1, S-s-1}$$

Convolution using Grouped Convolution

$$C_g = \frac{C}{G}$$

$$K_g = \frac{K}{G}$$

$$y_{n, k, p, q} = \sum_c^{C_g} \sum_r^R \sum_s^S x_{n, C_g * \text{floor}(k/K_g) + c, p+r, q+s} \times w_{k, c, r, s}$$

7.2. Grouped Convolutions

cuDNN supports grouped convolutions by setting `groupCount > 1` for the convolution descriptor `convDesc`, using `cudaSetConvolutionGroupCount()`.



Note: By default, the convolution descriptor `convDesc` is set to `groupCount` of 1.

Basic Idea

Conceptually, in grouped convolutions, the input channels and the filter channels are split into a `groupCount` number of independent groups, with each group having a reduced number of channels. The convolution operation is then performed separately on these input and filter groups.

For example, consider the following: if the number of input channels is 4, and the number of filter channels of 12. For a normal, ungrouped convolution, the number of computation operations performed are 12×4 .

If the `groupCount` is set to 2, then there are now two input channel groups of two input channels each, and two filter channel groups of six filter channels each.

As a result, each grouped convolution will now perform 2×6 computation operations, and two such grouped convolutions are performed. Hence the computation savings are 2x: $(12 \times 4) / (2 \times (2 \times 6))$.

cuDNN Grouped Convolution

- ▶ When using `groupCount` for grouped convolutions, you must still define all tensor descriptors so that they describe the size of the entire convolution, instead of specifying the sizes per group.
- ▶ Grouped convolutions are supported for all formats that are currently supported by the functions `cudaConvolutionForward()`, `cudaConvolutionBackwardData()` and `cudaConvolutionBackwardFilter()`.
- ▶ The tensor stridings that are set for `groupCount` of 1 are also valid for any group count.

- By default, the convolution descriptor `convDesc` is set to `groupCount` of 1.



Note: See [Convolution Formulas](#) for the math behind the cuDNN grouped convolution.

Example

Below is an example showing the dimensions and strides for grouped convolutions for NCHW format, for 2D convolution.



Note: The symbols `*` and `/` are used to indicate multiplication and division.

`xDesc` or `dxDesc`:

- Dimensions: `[batch_size, input_channel, x_height, x_width]`
- Strides: `[input_channels*x_height*x_width, x_height*x_width, x_width, 1]`

`wDesc` or `dwDesc`:

- Dimensions: `[output_channels, input_channels/groupCount, w_height, w_width]`
- Format: NCHW

`convDesc`:

- Group Count: `groupCount`

`yDesc` or `dyDesc`:

- Dimensions: `[batch_size, output_channels, y_height, y_width]`
- Strides: `[output_channels*y_height*y_width, y_height*y_width, y_width, 1]`

7.3. Best Practices For 3D Convolutions



ATTENTION: These guidelines are applicable to 3D convolution and deconvolution functions starting in CUDA® Deep Neural Network library™ (cuDNN) v7.6.3.

The following guidelines are for setting the cuDNN library parameters to enhance the performance of 3D convolutions. Specifically, these guidelines are focused on settings such as filter sizes, padding and dilation settings. Additionally, an application-specific use-case, namely, medical imaging, is presented to demonstrate the performance enhancement of 3D convolutions with these recommended settings.

Specifically, these guidelines are applicable to the following functions and their associated data types:

- [cudnnConvolutionForward\(\)](#)

- ▶ [cudnnConvolutionBackwardData\(\)](#)
- ▶ [cudnnConvolutionBackwardFilter\(\)](#)

For more information, refer to the [NVIDIA cuDNN Developer Guide](#) and the [NVIDIA cuDNN API Reference](#).

7.3.1. Recommended Settings

The following table shows the recommended settings while performing 3D convolutions for cuDNN.

Table 3. Recommended settings while performing 3D convolutions for cuDNN

		cuDNN 8.7.0
Platform		NVIDIA Hopper Architecture NVIDIA Ampere Architecture NVIDIA Turing Architecture NVIDIA Volta Architecture
Convolution (3D or 2D)		3D and 2D
Convolution or deconvolution (fprop, dgrad, or wgrad)		fprop dgrad wgrad
Grouped convolution size		C_per_group == K_per_group == {1, 4, 8, 16, 32, 64, 128, 256} Not supported for INT8
Data layout format (NHWC/NCHW) ¹		NDHWC
Input/output precision (FP16, FP32, INT8, or FP64)		FP16, FP32 ² , INT8 ³
Accumulator (compute) precision (FP16, FP32, INT32 or FP64)		FP32, INT32
Filter (kernel) sizes		No limitation
Padding		No limitation
Image sizes		2 GB limitation for a tensor
Number of channels	C	0 mod 8 0 mod 16 (for INT8)
	K	0 mod 8

¹ NHWC/NCHW corresponds to NDHWC/NCDHW in 3D convolution.

² With CUDNN_TENSOROP_MATH_ALLOW_CONVERSION pre-Ampere. Default TF32 math in NVIDIA Ampere Architecture.

³ INT8 does not support dgrad and wgrad. INT8 3D convolutions are only supported in the backend API. Refer to the tables in the [cudnnConvolutionForward\(\)](#) section for more information.

		cuDNN 8.7.0
		0 mod 16 (for INT8)
Convolution mode		Cross-correlation and convolution
Strides		No limitation
Dilation		No limitation
Data pointer alignment		All data pointers are 16-bytes aligned.

7.3.2. Limitations

Your application will be functional but could be less performant if the model has channel counts lower than 32 (gets worse the lower it is).

If the above is in the network, use `cuDNNFind` to get the best option.

Chapter 8. Features Of RNN Functions

Refer to the table below for a list of features supported by each RNN function:



Note:

For each of these terms, the short-form versions shown in the parenthesis are used in the tables below for brevity: `CUDNN_RNN_ALGO_STANDARD` (`_ALGO_STANDARD`), `CUDNN_RNN_ALGO_PERSIST_STATIC` (`_ALGO_PERSIST_STATIC`), `CUDNN_RNN_ALGO_PERSIST_DYNAMIC` (`_ALGO_PERSIST_DYNAMIC`), and `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` (`_ALLOW_CONVERSION`).

Functions	Input/output layout supported	Supports variable sequence length in batch	Commonly supported
<code>cudaRnnForwardInferOnly()</code> <code>cudaRnnForwardTraining()</code> <code>cudaRnnBackwardData()</code> <code>cudaRnnBackwardWeights()</code>	Only Sequence major, packed (non-padded)	Only with <code>_ALGO_STANDARD</code> Require input sequences descending sorted according to length.	Mode (cell type) supported: <code>CUDNN_RNN_RELU</code> , <code>CUDNN_RNN_TANH</code> , <code>CUDNN_LSTM</code> , <code>CUDNN_GRU</code> Algo supported ⁴ (see the table below for an elaboration on these algorithms): <code>_ALGO_STANDARD</code> , <code>_ALGO_PERSIST_STATIC</code> , <code>_ALGO_PERSIST_DYNAMIC</code>
<code>cudaRnnForwardInferOnlyEx()</code> <code>cudaRnnForwardTrainingEx()</code> <code>cudaRnnBackwardDataEx()</code> <code>cudaRnnBackwardWeightsEx()</code>	Sequence major unpacked Batch major unpacked ⁵ Sequence major packed ⁶	Only with <code>_ALGO_STANDARD</code> For unpacked layout, no input sorting required. ⁷ For packed layout, require input sequences descending sorted according to length.	Math mode supported: <code>CUDNN_DEFAULT_MATH</code> , <code>CUDNN_TENSOR_OP_MATH</code> (will automatically fall back if run on pre-Volta or if algo doesn't support Tensor Cores)

⁴ Do not mix different algos for different steps of training. It's also not recommended to mix non-extended and extended API for different steps of training.

⁵ To use an unpacked layout, users need to set `CUDNN_RNN_PADDED_IO_ENABLED` through `cudaRnnSetRnnPaddingMode()`.

⁷ To use an unpacked layout, users need to set `CUDNN_RNN_PADDED_IO_ENABLED` through `cudaRnnSetRnnPaddingMode()`.

⁶ To use an unpacked layout, users need to set `CUDNN_RNN_PADDED_IO_ENABLED` through `cudaRnnSetRnnPaddingMode()`.

Functions	Input/output layout supported	Supports variable sequence length in batch	Commonly supported
			_ALLOW_CONVERSION (may do down conversion to utilize Tensor Cores) Direction mode supported: CUDNN_UNIDIRECTIONAL, CUDNN_BIDIRECTIONAL RNN input mode: CUDNN_LINEAR_INPUT, CUDNN_SKIP_INPUT

The following table provides the features supported by the algorithms referred in the above table: CUDNN_RNN_ALGO_STANDARD, CUDNN_RNN_ALGO_PERSIST_STATIC, and CUDNN_RNN_ALGO_PERSIST_DYNAMIC.

Features	_ALGO_STANDARD	_ALGO_PERSIST_STATIC	CUDNN_RNN_ALGO_PERSIST_DYNAMIC	_ALGO_PERSIST_DYNAMIC	DSO_MATCH
Half input	Supported				
Single accumulation	Half intermediate storage				
Half output	Single accumulation				
Single input	Supported				
Single accumulation	If running on Volta, with CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION*, will down-convert and use half intermediate storage.				
Single output	Otherwise: Single intermediate storage				
	Single accumulation				
Double input	Supported	Not Supported	Not Supported	Supported	
Double accumulation	Double intermediate storage			Double intermediate storage	
Double output	Double accumulation			Double accumulation	
LSTM recurrent projection	Supported	Not Supported	Not Supported	Not Supported	
LSTM cell clipping	Supported				
Variable sequence length in batch	Supported	Not Supported	Not Supported	Not Supported	
Tensor Cores	Supported			Not Supported, will execute normally ignoring	
	For half input/output, acceleration requires setting				

Features	<code>_ALGO_STANDARD</code>	<code>_ALGO_PERSISTENT</code>	<code>CUDNN_RNN_ALGO_STANDARD</code>	<code>_ALGO_PERSISTENT_DYNAMIC</code>
	<p><code>CUDNN_TENSOR_OP_MATH</code>⁸ or <code>CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION</code>⁹</p> <p>Acceleration requires <code>inputSize</code> and <code>hiddenSize</code> to be a multiple of 8</p> <p>For single input/output on NVIDIA Volta, NVIDIA Xavier, and NVIDIA Turing, acceleration requires setting <code>CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION</code>¹⁰</p> <p>Acceleration requires <code>inputSize</code> and <code>hiddenSize</code> to be a multiple of 8</p> <p>For single input/output on NVIDIA Ampere Architecture, acceleration requires setting <code>CUDNN_DEFAULT_MATH, CUDNN_TENSOR_OP_MATH, Or CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION*</code></p> <p>Acceleration requires <code>inputSize</code> and <code>hiddenSize</code> to be a multiple of 4</p>			<p><code>CUDNN_TENSOR_OP_MATH</code>¹¹ or <code>_ALLOW_CONVERSION</code>¹²</p>
Other limitations		Max problem size is limited by GPU specifications.	<p>Forward RNN:</p> <ul style="list-style-type: none"> ▶ RELU and TANH RNN: <code>hidden_size</code> <= 384 ▶ LSTM and GRU: <code>hidden_size</code> <= 192 <p>BackwardData RNN:</p> <ul style="list-style-type: none"> ▶ RELU and TANH RNN: <code>hidden_size</code> <= 256 ▶ LSTM and GRU: <code>hidden_size</code> <= 128 	Requires real time compilation through NVRTC

¹¹ `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` can be set through `cudaSetRNNMatrixMathType()`.

⁸ `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` can be set through `cudaSetRNNMatrixMathType()`.

⁹ `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` can be set through `cudaSetRNNMatrixMathType()`.

¹² `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` can be set through `cudaSetRNNMatrixMathType()`.

¹⁰ `CUDNN_TENSOR_OP_MATH` or `CUDNN_TENSOR_OP_MATH_ALLOW_CONVERSION` can be set through `cudaSetRNNMatrixMathType()`.

Chapter 9. Mixed Precision Numerical Accuracy

When the computation precision and the output precision are not the same, it is possible that the numerical accuracy will vary from one algorithm to the other.

For example, when the computation is performed in FP32 and the output is in FP16, the `CUDNN_CONVOLUTION_BWD_FILTER_ALGO_0` (`ALGO_0`) has lower accuracy compared to the `CUDNN_CONVOLUTION_BWD_FILTER_ALGO_1` (`ALGO_1`). This is because `ALGO_0` does not use extra workspace, and is forced to accumulate the intermediate results in FP16, i.e., half precision float, and this reduces the accuracy. The `ALGO_1`, on the other hand, uses additional workspace to accumulate the intermediate values in FP32, i.e., full precision float.

Chapter 10. The cuDNN Graph API

The cuDNN library provides a declarative programming model for describing computation as a graph of operations. This *graph API* was introduced in cuDNN 8.0 to provide a more flexible API, especially with the growing importance of operation fusion.

The user starts by building a graph of operations, like the one pictured in [#unique_52/unique_52_Connect_42_fig_w12_btl_vlb](#).

At a high level, the user is describing a dataflow graph of operations on tensors. Given a *finalized* graph, the user then selects and configures an engine that can execute that graph. There are several methods for selecting and configuring engines, which have tradeoffs with respect to ease-of-use, runtime overhead, and engine performance. The next subsection walks through an example operation graph, covering the process in more detail.

The graph API has two entry points:

- ▶ [C backend API](#) (lowest level entry point into the graph API)
- ▶ [C++ frontend API](#) (convenience layer on top of the C backend API)

We expect that most users prefer the C++ frontend API because:

- ▶ It is less verbose without loss of control - all functionality accessible through the backend API is also accessible through the frontend API.
- ▶ It adds functionality on top of the backend API, like errata filters and autotuning.
- ▶ It is open source.

In either case (i.e. the backend or frontend API), the high level concepts are the same.

10.1. Graph API Example with Operation Fusion

In the following example, the user would like to implement a fusion operation of convolution, bias, and activation.

10.1.1. Creating Operation and Tensor Descriptors to Specify the Graph Dataflow

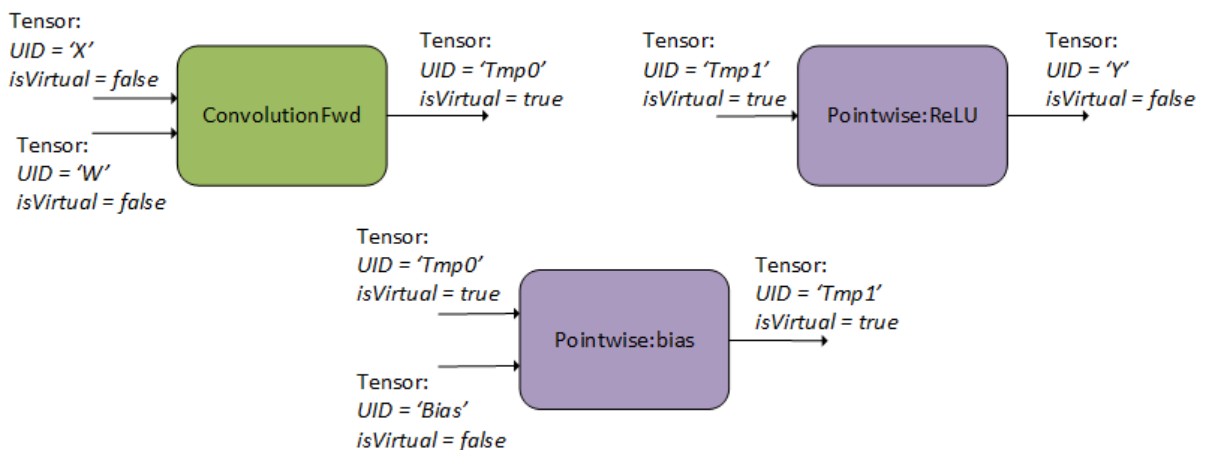
First, create three cuDNN backend operation descriptors.

As can be seen in [Figure 9](#), the user specified one forward convolution operation (using `CUDNN_BACKEND_OPERATION_CONVOLUTION_FORWARD_DESCRIPTOR`), a pointwise operation for the bias addition (using `CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR` with mode `CUDNN_POINTWISE_ADD`), and a pointwise operation for the ReLU activation (using `CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR` with mode `CUDNN_POINTWISE_RELU_FWD`). Refer to the [backend API](#) for more details on setting the attributes of these descriptors. For an example of how a forward convolution can be set up, refer to the [use case](#) in the backend API.

The user should also create tensor descriptors for the inputs and outputs of all of the operations in the graph. The graph dataflow is implied by the assignment of tensors (refer to [Figure 9](#)), for example, by specifying the backend tensor *Tmp0* as both the output of the convolution operation and the input of the bias operation, cuDNN infers that the dataflow runs from the convolution into the bias. The same applies to tensor *Tmp1*. If the user doesn't need the intermediate results *Tmp0* and *Tmp1* for any other use, then the user can specify them to be virtual tensors, so the memory I/Os can later be optimized out.

- Note that graphs with more than one operation node do not support in-place operations (that is, where any of the input UIDs matches any of the output UIDs). Such in-place operations are considered cyclic in later graph analysis and deemed unsupported. In-place operations are supported for single-node graphs.
- Also note that the operation descriptors can be created and passed into cuDNN in any order, as the tensor UIDs are enough to determine the dependencies in the graph.

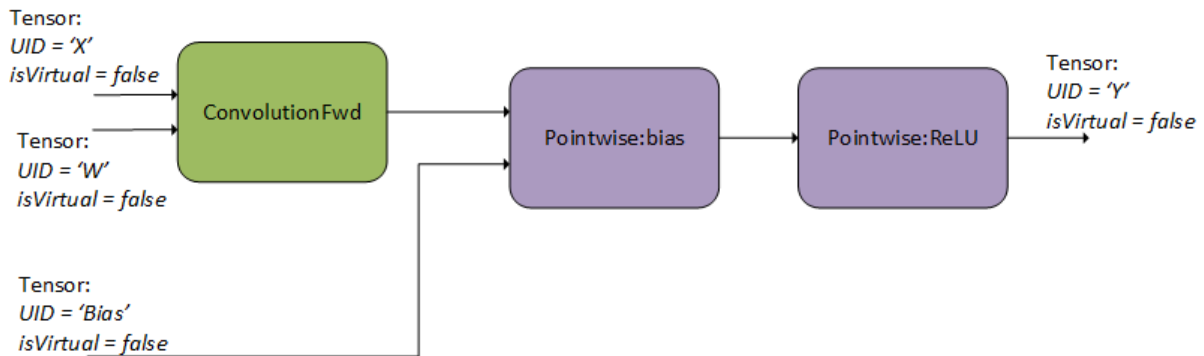
Figure 9. A set of operation descriptors the user passes to the operation graph



10.1.2. Finalizing The Operation Graph

Second, the user finalizes the operation graph. As part of finalization, cuDNN performs the dataflow analysis to establish the dependency relationship between operations and connect the edges, as illustrated in the following figure. In this step, cuDNN performs various checks to confirm the validity of the graph.

Figure 10. The operation graph after finalization



10.1.3. Configuring An Engine That Can Execute The Operation Graph

Third, given the finalized operation graph, the user must select and configure an engine to execute that graph, which results in an execution plan. There are three methods:

1. **Heuristics.** Users that prefer cuDNN to recommend the best engine and knob choices can query cuDNN's heuristics to get a list of engine configs, sorted by predicted performance. Typically, the user constructs the execution plan using the top ranked engine config in the list.
2. **Auto-tuning.** The user can also iterate over a list of engine configs and time each one to choose the best engine config for a particular problem on a particular device. The C++ frontend API provides a convenience function, `cudaDnnFindPlan`, which does this. To reduce overhead, a user might only auto-tune over the top N engine configs returned by the heuristics.
3. **Manual.** Expert users can query for all engines that can support the operation graph. For each engine, the user can then further query the numerical notes and adjustable knobs. Numerical notes inform the user about the numerical behavior of the engine such as whether it does datatype down conversion at the input or during output reduction. The adjustable knobs allow fine grained control of the engine's behavior and performance. With the engine choice and the knob choice determined, the user can construct the engine, engine config, and execution plan.

10.1.4. Executing The Engine

Finally, with the execution plan constructed and when it comes time to run it, the user should construct the backend variant pack by providing the workspace pointer, an array of UIDs, and an array of device pointers. The UIDs and the pointers should be in the corresponding order. With the handle, the execution plan and variant pack, the execution API can be called and the computation is carried out on the GPU.

10.2. Supported Graph Patterns

The cuDNN Graph API supports a set of graph patterns. These patterns are supported by a large number of engines, each with their own support surfaces. These engines are grouped into three different classes, as reflected by the following three subsections: pre-compiled single operation engines, runtime fusion engines, and specialized pre-compiled engines.

Since these engines have some overlap in the patterns they support, a given pattern may result in zero, one, or more engines.


10.2.1. Pre-compiled Single Operation Engines

One basic class of engines includes pre-compiled engines that support an operation graph with just one operation; specifically: `ConvolutionFwd`, `ConvolutionBwFilter`, `ConvolutionBwData`, or `ConvolutionBwBias`. Their more precise support surface can be found in the [NVIDIA cuDNN API Reference](#).

10.2.1.1. ConvolutionFwd

`ConvolutionFwd` computes the convolution of X with filter data W . In addition, it uses scaling factors α and β to blend this result with the previous output. This graph operation is similar to [`cudaConvolutionForward\(\)`](#).


Figure 11. ConvolutionFwd Engine



10.2.1.2. ConvolutionBwFilter

`ConvolutionBwFilter` computes the convolution filter gradient of the tensor dy . In addition, it uses scaling factors α and β to blend this result with the previous output. This graph operation is similar to [`cudaConvolutionBackwardFilter\(\)`](#).


Figure 12. ConvolutionBwFilter Engine



10.2.1.3. ConvolutionBwData

`ConvolutionBwData` computes the convolution data gradient of the tensor `dy`. In addition, it uses scaling factors α and β to blend this result with the previous output. This graph operation is similar to `cudnnConvolutionBackwardData()`.

Figure 13. ConvolutionBwData Engine



10.2.2. Runtime Fusion Engine

The engines documented in the previous section support single-op patterns. Of course, for fusion to be interesting, the graph needs to support multiple operations. And ideally, we want the supported patterns to be flexible to cover a diverse set of use cases. To accomplish this generality, cuDNN has a runtime fusion engine that generates the kernel (or kernels) at runtime based on the graph pattern. This section outlines the patterns supported by these runtime fusion engines (that is, engines with `CUDNN_BEHAVIOR_NOTE_RUNTIME_COMPILATION` behavioral note).

We can think of the support surface as covering the following generic patterns:

1. `ConvolutionFwd` fusions

$$g_2(Y = \text{convolutionFwd}(X = g_1(\text{inputs}), W), \text{inputs})$$

2. `ConvolutionBwFilter` fusions

$$g_2(dw = \text{convolutionBwFilter}(dy = g_1(\text{inputs}), X), \text{inputs})$$

3. `ConvolutionBwData` fusions

$$g_2(dx = \text{convolutionBwData}(dy, W), \text{inputs})$$

4. `MatMul` fusions

$$g_2(C = \text{matmul}(A = g_1(\text{inputs}), B), \text{inputs})$$

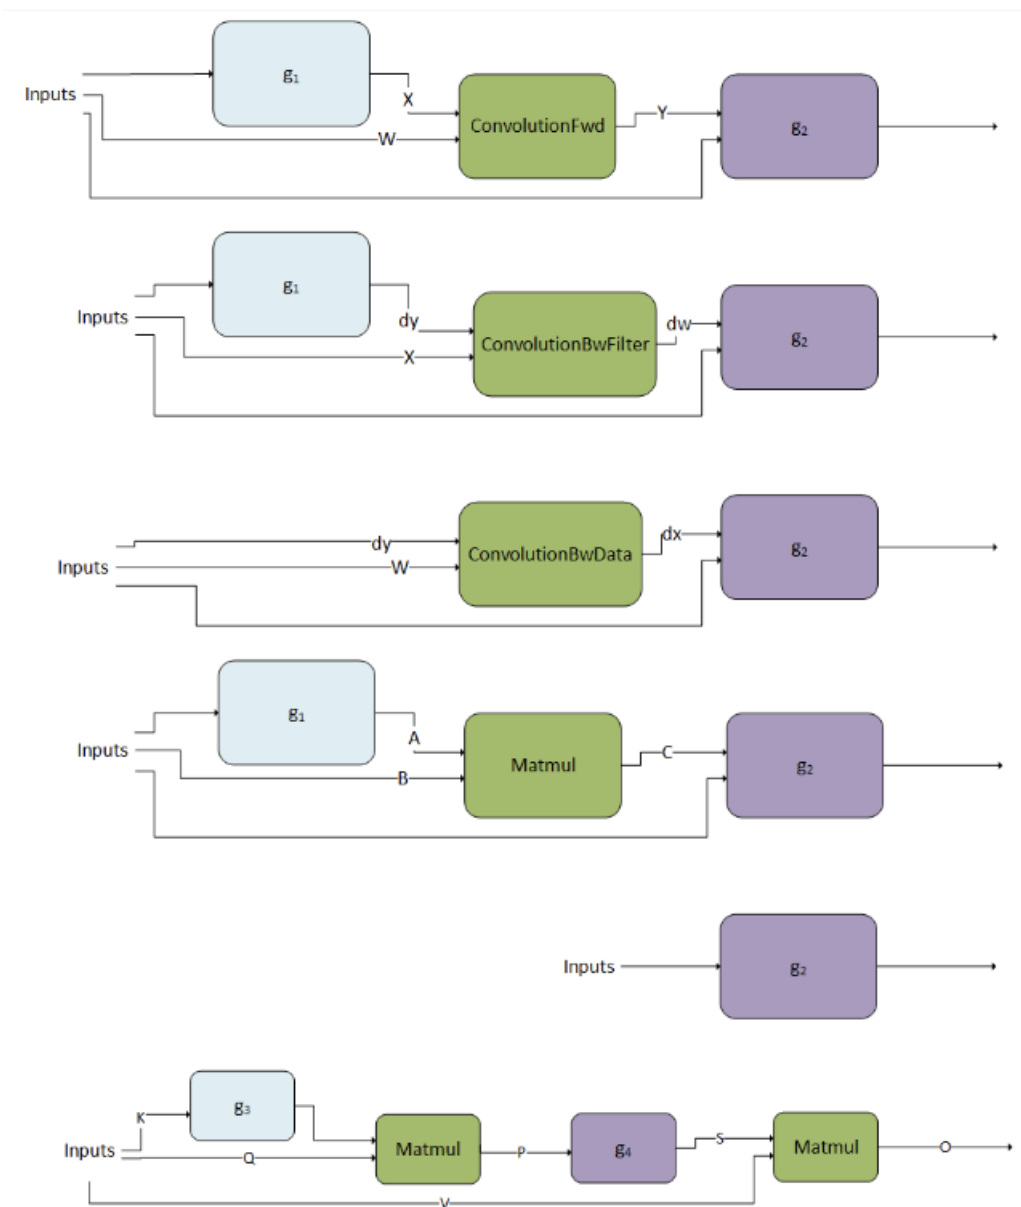
5. Pointwise fusions

$$g_2(\text{inputs})$$

6. MatMul-MatMul fusions

$$O = \text{matmul}(g_4(P = \text{matmul}(Q, g_2(K)), V))$$

Figure 14. Graphical Representation of the following Generic Patterns Supported by the Runtime Fusion Engine



g_1 is a directed acyclic graph (DAG) that can consist of zero or any number of the following operation:

- ▶ CUDNN_BACKEND_OPERATION_CONCAT_DESCRIPTOR
- ▶ CUDNN_BACKEND_OPERATION_SIGNAL_DESCRIPTOR
- ▶ CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR

g_2 is a DAG that can consist of zero or any number of the following operations:


- ▶ CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR
- ▶ CUDNN_BACKEND_OPERATION_RESAMPLE_FWD_DESCRIPTOR
- ▶ CUDNN_BACKEND_OPERATION_RESAMPLE_BWD_DESCRIPTOR
- ▶ CUDNN_BACKEND_OPERATION_GEN_STATS_DESCRIPTOR
- ▶ CUDNN_BACKEND_OPERATION_REDUCTION_DESCRIPTOR
- ▶ CUDNN_BACKEND_OPERATION_SIGNAL_DESCRIPTOR

MatMul-MatMul has been added to the runtime fusion engine to serve patterns that are commonly used in multihead attention. g_3 and g_4 are limited to a certain few DAGs of operations.

g_3 can be an empty graph or a single scale operation with the scale being a scalar value (CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR with mode CUDNN_POINTWISE_MUL).

g_4 can be empty or the combination of the following DAGs of cuDNN operations. Each of these DAGs is optional, as shown by the dotted line.

Figure 15. DAGs of cuDNN operations




The combination has to obey the order in which we present them. For example: if you want to use the padding mask and softmax, the padding mask has to appear before softmax.

These operations are commonly used in multihead attention. In the following diagram, we depict how to create a DAG for each of the operations. In later versions, we will be expanding the possible DAGs for g_3 and g_4 .


Padding Mask

Figure 16. cuDNN graph depicting DAG:Padding Mask




Softmax

Figure 17. cuDNN graph depicting DAG:Softmax



Dropout

Figure 18. cuDNN graph depicting DAG:Dropout



Note:

- ▶ The arrow going into g_2 can go into any of g_2 's nodes and does not necessarily need to feed into a root node.
- ▶ The abbreviated notations for operations are used in the diagrams and throughout the text for visualization purposes.

10.2.2.1. Limitations

While the generic patterns listed previously are widely applicable, there are some cases where we do not have full support.

Limitations Common to all Generic Patterns

Limitations to g_1 :

- ▶ Concatenation or signaling operations, if present, should be before any pointwise operations.
- ▶ For compute capability < 8.0, g_1 is not supported.

Limitations to g_2 :

- ▶ As specified in the previous section, g_2 can include only `Pointwise` operations, `ResampleFwd`, `ResampleBwd`, `GenStats`, and `Reduction`.
- ▶ The I/O (that is, non-virtual) tensor data type can be any of `{FP32, FP16, BF16, INT8, packed-BOOLEAN}`.
- ▶ For pointwise operations, non-virtual tensors need to be either all NCHW (or row-major), or all NHWC (or column-major).
- ▶ The intermediate virtual tensor data type can be any of `{FP32, FP16, BF16, INT8, BOOLEAN}`, and this intermediate storage type is obeyed by the code-generator. Generally, FP32 is recommended.
- ▶ The input tensor to a `ResampleFwd` or `ResampleBwd` operation should not be produced by another operation within this graph, but should come from global memory. The two operations cannot be used in the `ConvolutionBwFilter`, `ConvolutionBwData`, and `MatMul` fusion patterns.
- ▶ There can be at most one reduction operation, and it needs to be at the final node of g_2 .
- ▶ Signaling operations, if present, must be the final nodes in g_2 . Hence, signaling operations cannot be used in conjunction with reduction operations.

Limitations per Generic Pattern

Table 4. Limitations per Generic Pattern

	Limitations to g_1
ConvolutionFwd fusions	<ul style="list-style-type: none"> ▶ Fusion operations on input tensors can be only a chain of three specific pointwise operations, in this exact order: <code>Pointwise:mul</code>, <code>Pointwise:add</code>, and <code>Pointwise:ReLU</code>. This specific support is

	Limitations to g_1
	<p>added to realize convolution batch norm fusion use cases.</p> <ul style="list-style-type: none"> ▶ All tensors involved can only be FP16. ▶ <code>Pointwise:mul</code> can only be with a tensor of scalars per channel. ▶ <code>Pointwise:add</code> can only be a column broadcast.
ConvolutionBwFilter fusions	Same limitations specified for ConvolutionFwd fusions apply here.
ConvolutionBwData fusions	No fusion on input tensors for backward data convolution is supported.
MatMul fusions	<ul style="list-style-type: none"> ▶ Can be any combination of pointwise operations. ▶ Only fusible with operand A, not with B. ▶ Operand A should have an FP16 data type. ▶ Broadcasted input can have any data type. ▶ Compute type is FP32 only.
Pointwise fusions	Not Applicable

Table 5. Limitations per Generic Pattern

	Limitations to Matmul-Matmul fusions
MatMul	<ul style="list-style-type: none"> ▶ Compute type for both MatMul ops needs to be float ▶ Input tensors need to have datatype FP16 or BF16 ▶ Output tensors need to have datatype FP16, BF16, or FP32 (TF32) datatype
Pointwise operations in g_3 and g_4	Compute type needs to be FP32 (TF32)
Reduction operations in g_3 and g_4	I/O types and compute type needs to be FP32 (TF32)
RNG operation in g_3 and g_4	<ul style="list-style-type: none"> ▶ Data type of <code>yTensor</code> needs to be FP32 (TF32) ▶ The <code>CUDNN_TYPE_RNG_DISTRIBUTION</code> needs to be <code>CUDNN_RNG_DISTRIBUTION_BERNOULLI</code>

Tensor Layout Requirements

Lastly, there are some layout requirements to the I/O tensors involved in fusion graphs. For more information, refer to the [Tensor Descriptor](#) and [Data Layout Formats](#) sections. The following table describes the requirements per fusion pattern:

Table 6. Layout Requirements per Pattern

Pattern	Layout Requirement
ConvolutionFwd, ConvolutionBwFilter, ConvolutionBwData fusions	<ul style="list-style-type: none"> ▶ All tensors are fully packed NHWC.
MatMul fusions	<ul style="list-style-type: none"> ▶ Input operands can have either row-major or all column-major. ▶ In g_1, the tensor operating with Matrix A (dim[B, M, K]) can be either a scalar with dim[1, 1, 1], a row vector with dim[B, M, 1], a column vector with dim[B, 1, K], or a full matrix with dim[B, M, K]. ▶ In g_2, all I/O tensors should be either all row-major or all column-major.
Pointwise fusions	<ul style="list-style-type: none"> ▶ If all tensors are 3D, the same layout requirements as MatMul g_2. ▶ If all tensors are 4D or 5D, the same requirements as ConvolutionFwd, ConvolutionBwFilter, ConvolutionBwData layout.
Matmul-Matmul fusions	<ul style="list-style-type: none"> ▶ All I/O tensors need to have 4 dimensions, with the first two denoting the batch dimensions. The usage of rank-4 tensors in MatMul ops can be read from the backend API documentation. ▶ The contracting dimension (dimension K) for the first matmul needs to be 64. ▶ The non contracting dimension (dimensions M and N) for the first matmul needs to be equal and less than or equal to 512. ▶ The packing of the tensors for Q, V and O tensor is expected to follow the pattern:

Pattern	Layout Requirement
	<p>for dimensions equal to [d1, d2, d3, d4]; the strides are expected to be [d2 * d3 * d4, d4, d2 * d4, 1]</p> <p>► For K tensor: for dimensions equal to [d1, d2, d4, d3]; the strides are expected to be [d2 * d3 * d4, d4, 1, d2 * d4]</p>

10.2.2.2. Examples of Supported Patterns


The following sections provide examples of supported patterns, in order of increasing complexity. We employ the same color scheme as in the overall pattern to aid in identifying the structure of g_1 (blue) and g_2 (purple).

For illustration purposes, we abbreviated the operations used. For a full mapping to the actual backend descriptors, refer to the [Mapping with Backend Descriptors](#).

10.2.2.2.1. Single Operation

The following example illustrates a convolution operation without any operations before or after it. This means, g_1 and g_2 , are empty graphs.


Figure 19. This example illustrates the Runtime Fusion Engine with a Single Operation



10.2.2.2.2. Pointwise Operations After Convolution 1

In this example, g_2 consists of a sequential set of two pointwise operations after the convolution.


Figure 20. ConvolutionFwd Followed by a DAG with Two Operations



10.2.2.2.3. Pointwise Operations After Convolution 2

Similar to the previous example, g_2 consists of a sequential set of multiple pointwise operations.


Figure 21. ConvolutionFwd Followed by a DAG with Three Operations



10.2.2.2.4. Pointwise Operations Before Matrix Multiplication

Pointwise operations can also precede a convolution or matrix multiplication, that is, g_1 is composed of pointwise operations.

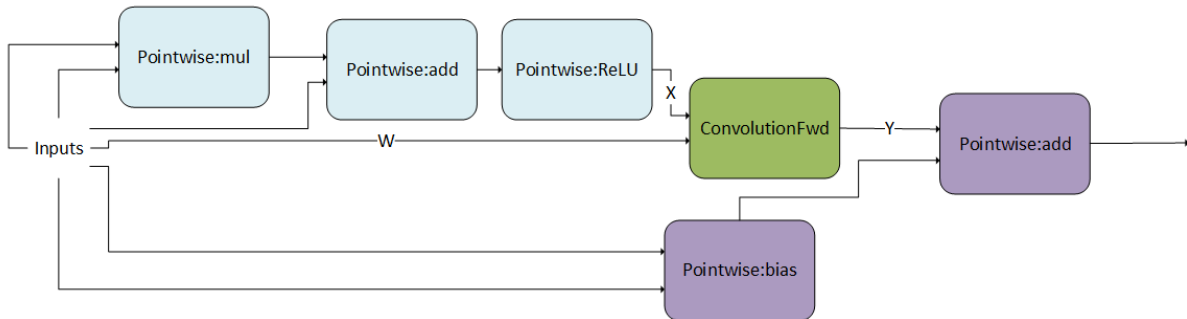
Figure 22. MatMul Preceded by a DAG with Two Operations



10.2.2.2.5. Convolution Producer Node in Middle of DAG

The following pattern shows g_1 as a DAG of pointwise operations feeding into a convolution. In addition, g_2 is a DAG consisting of two pointwise operations. Note that the convolution is being consumed in the middle of g_2 as opposed to g_2 's first node. This is a valid pattern.

Figure 23. This example illustrates fusion of operations before and after the `ConvolutionFwd` operation. In addition we observe that the output of `ConvolutionFwd` can feed anywhere in g_2 .



10.2.2.3. Operation specific Constraints for the Runtime Fusion Engine

Every operation in the supported generic patterns of the runtime fusion engine is subject to a few specific constraints regarding their parameter surface. The following subsections document these.

Note that these constraints are in addition to (1) any constraints mentioned in the [Backend API](#), and (2) limitations in relation to other operations in the directed acyclic graph (DAG), as mentioned in the [Limitations](#) section.

10.2.2.3.1. Convolutions

There are three operation nodes that represent different types of convolutions namely:

ConvolutionFwd

This operation represents forward convolution, that is, computing the response tensor of image tensor convoluted with filter tensor. For complete details on the interface, as well as general constraints, refer to the [CUDNN_BACKEND_OPERATION_CONVOLUTION_FORWARD_DESCRIPTOR](#) section.

ConvolutionBwFilter

This operation represents convolution backward filters, that is, computing filter gradients from a response and an image tensor. For complete details on the interface, as well as general constraints, refer to the [CUDNN_BACKEND_OPERATION_CONVOLUTION_BACKWARD_FILTER_DESCRIPTOR](#) section.

ConvolutionBwData

This operation represents convolution backward data, that is, computing input data gradients from a response and a filter tensor. For complete details on the interface, as well as general constraints, refer to the [CUDNN_BACKEND_OPERATION_CONVOLUTION_BACKWARD_DATA_DESCRIPTOR](#) section.

Table 7. Tensor Attributes for all Three Operations

	Input Tensor Attribute Name	Output Tensor Attribute Name
ConvolutionFwd	CUDNN_ATTR_OPERATION_CONVOLUTION_FORWARD_X CUDNN_ATTR_OPERATION_CONVOLUTION_FORWARD_W	CUDNN_ATTR_OPERATION_CONVOLUTION_FORWARD_X CUDNN_ATTR_OPERATION_CONVOLUTION_FORWARD_W
ConvolutionBwFilter	CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_DATA_DX CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_DATA_DY	CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_DATA_DX CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_DATA_DY
ConvolutionBwData	CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_FILTER_DX CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_FILTER_DY	CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_FILTER_DX CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_FILTER_DY

The following tables list the constraints for all three operations, in addition to any constraints mentioned in the [Backend API](#), and any constraints listed in the [Limitations](#) section, in relation to other operations. Note that these additional constraints only apply when these operations are used in the runtime fusion engine.

Table 8. Constraints for all Three Operations

Attribute	Support
CUDNN_ATTR_CONVOLUTION_MODE	CUDNN_CROSS_CORRELATION
CUDNN_ATTR_CONVOLUTION_COMP_TYPE	<ul style="list-style-type: none"> ► For ConvolutionFwd CUDNN_DATA_HALF, CUDNN_DATA_INT32, and CUDNN_DATA_FLOAT ► For ConvolutionBwData and ConvolutionBwFilter ► Only CUDNN_DATA_FLOAT
CUDNN_ATTR_CONVOLUTION_SPATIAL_DIMS	2 or 3
CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_FILTER_ALPHA	0.1 to 1.0
CUDNN_ATTR_OPERATION_CONVOLUTION_BWD_FILTER_BETA	0.1 to 1.0

Table 9. I/O Tensors Alignment Requirements

Tensor Data Type	Number of input and output channels for NVIDIA Hopper Architecture	Number of input and output channels for NVIDIA Ampere and Ada Lovelace	Number of input and output channels for NVIDIA Volta/Turing Architecture
INT8	Multiple of 4	Multiple of 4	Multiple of 16
FP8	Multiple of 16	N/A	N/A
FP16/BF16	Multiple of 2	Multiple of 2	Multiple of 8
FP32 (TF32)	Any value	Any value	Multiple of 4

Lastly, there are some batch size requirements per operation:

Table 10. Batch Size Requirements Per Operation

Operation	Batch size for FP8 data type on NVIDIA Hopper Architecture	Batch size for other data types
ConvolutionFwd	Any	Any
ConvolutionBwFilter	Multiple of 16	Any
ConvolutionBwData	Multiple of 16	Any

The FP8 data type since Hopper architecture has two variants; `CUDNN_DATA_FP8_E4M3` and `CUDNN_DATA_FP8_E5M2` as I/O data types. It also has two possible compute types; `CUDNN_DATA_FLOAT` and `CUDNN_DATA_FAST_FLOAT_FOR_FP8`, which is a faster, but less accurate option for FP8 Tensor Core operations. It is sufficiently accurate for inference or the forward pass of training. However, for FP8 training backward pass computations (that is, computing weight and activation gradients), we recommend choosing the more accurate `CUDNN_DATA_FLOAT` compute type to preserve a higher level of accuracy which can be necessary for some models.

Table 11. Recommended compute type for FP8 tensor computations for Hopper architecture

Operation	Recommended I/O type	Recommended compute type
ConvolutionFwd	<code>CUDNN_DATA_FP8_E4M3</code>	<ul style="list-style-type: none"> ▶ <code>CUDNN_DATA_FAST_FLOAT_FOR_FP8</code> ▶ <code>CUDNN_DATA_FLOAT</code>
ConvolutionBwData	<code>CUDNN_DATA_FP8_E4M3</code>	<code>CUDNN_DATA_FLOAT</code>
BatchNorm	<code>CUDNN_DATA_FP8_E4M3</code>	<code>CUDNN_DATA_FLOAT</code>
Pooling	<ul style="list-style-type: none"> ▶ <code>CUDNN_DATA_FP8_E4M3</code> ▶ <code>CUDNN_DATA_FP8_E5M2</code> 	<code>CUDNN_DATA_FLOAT</code>
Pointwise	<ul style="list-style-type: none"> ▶ <code>CUDNN_DATA_FP8_E4M3</code> ▶ <code>CUDNN_DATA_FP8_E5M2</code> 	<code>CUDNN_DATA_FLOAT</code>

10.2.2.3.2. MatMul

This operation represents matrix-matrix multiplication: $A * B = C$. For complete details on the interface, refer to the [CUDNN_BACKEND_OPERATION_MATMUL_DESCRIPTOR](#) section.

The following two tables list the constraints for MatMul operations, in addition to any general constraints as listed in the [Backend API](#), and any constraints listed in

the [Limitations](#) section, in relation to other operations. Note that these additional constraints only apply when MatMul is used in the runtime fusion engine.

Table 12. Constraints for MatMul Operations

Attribute	Support
CUDNN_ATTR_MATMUL_COMP_TYPE	CUDNN_DATA_HALF, CUDNN_DATA_INT32, and CUDNN_DATA_FLOAT

Table 13. MatMul Alignment Requirements

Tensor Data Type	Innermost dimension for NVIDIA Ampere Architecture and later	Innermost dimension for NVIDIA Volta/Turing Architecture
INT8	Multiple of 4	Multiple of 16
FP16/BF16	Multiple of 2	Multiple of 8
FP32 (TF32)	Any value	Multiple of 4

10.2.2.3.3. Pointwise

Represents a pointwise operation that implements the equation $Y = \text{op}(\alpha_1 * X)$ or $Y = \text{op}(\alpha_1 * X, \alpha_2 * B)$. Refer to the [CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR](#) and [CUDNN_BACKEND_POINTWISE_DESCRIPTOR](#) sections for more information and general constraints.

The following table lists the constraints for pointwise operations, in addition to the general constraints listed above, and any constraints listed in the [Limitations](#) section, in relation to other operations. Note that these additional constraints only apply when these operations are used in the runtime fusion engine.

Table 14. Constraints for Pointwise Operations

Attribute	Requirement
Tensor data type for CUDNN_ATTR_OPERATION_POINTWISE_XDESC, CUDNN_ATTR_OPERATION_POINTWISE_YDESC and, if applicable, CUDNN_ATTR_OPERATION_POINTWISE_BDESC	<ul style="list-style-type: none"> ▶ For any of the logical operators (CUDNN_POINTWISE_LOGICAL_AND, CUDNN_POINTWISE_LOGICAL_OR, and CUDNN_POINTWISE_LOGICAL_NOT), data type can be any of CUDNN_DATA_INT32, CUDNN_DATA_INT8, or CUDNN_DATA_BOOLEAN. ▶ For all other operators, all data types are supported.
CUDNN_ATTR_POINTWISE_MATH_PREC	<ul style="list-style-type: none"> ▶ For any of the logical operators (CUDNN_POINTWISE_LOGICAL_AND, CUDNN_POINTWISE_LOGICAL_OR, and

Attribute	Requirement
	<p>CUDNN_POINTWISE_LOGICAL_NOT), math precision needs to be CUDNN_DATA_BOOLEAN.</p> <ul style="list-style-type: none"> ► For all other operators, only CUDNN_DATA_FLOAT is supported.
CUDNN_ATTR_OPERATION_POINTWISE_ALPHA1	1.0f
CUDNN_ATTR_OPERATION_POINTWISE_ALPHA2	1.0f

10.2.2.3.4.GenStats

Represents an operation that generates per-channel statistics. Refer to the [CUDNN_BACKEND_OPERATION_GEN_STATS_DESCRIPTOR](#) section for more information and general constraints.

The following table lists the constraints for GenStats operations, in addition to the general constraints listed above, and any constraints listed in the [Limitations](#) section, in relation to other operations. Note that these additional constraints only apply when GenStats operations are used in the runtime fusion engine.

Table 15. Constraints for GenStats Operations

Attribute	Requirement
Tensor data type for CUDNN_ATTR_OPERATION_GENSTATS_XDESC	<ul style="list-style-type: none"> ► Prior to the NVIDIA Ampere Architecture GPU: CUDNN_DATA_HALF ► On NVIDIA Ampere Architecture and later: CUDNN_DATA_HALF and CUDNN_DATA_FLOAT
Tensor shape for CUDNN_ATTR_OPERATION_GENSTATS_SUMDESC and CUDNN_ATTR_OPERATION_GENSTATS_SQSUMDESC	Both should be of shape [1, C, 1, 1] for 2D conv or [1, C, 1, 1, 1] for 3D conv.
Tensor data type for CUDNN_ATTR_OPERATION_GENSTATS_SUMDESC and CUDNN_ATTR_OPERATION_GENSTATS_SQSUMDESC	CUDNN_DATA_FLOAT
CUDNN_ATTR_POINTWISE_MATH_PREC	CUDNN_DATA_FLOAT
Tensor layout for CUDNN_ATTR_OPERATION_GENSTATS_XDESC, CUDNN_ATTR_OPERATION_GENSTATS_SUMDESC and CUDNN_ATTR_OPERATION_GENSTATS_SQSUMDESC	NHWC fully packed

10.2.2.3.5.Reduction

This operation represents reducing values of a tensor in one or more dimensions. Refer to the [CUDNN_BACKEND_OPERATION_REDUCTION_DESCRIPTOR](#) section for more information and general constraints.

The following two tables are constraints for Reduction forward operations, in addition to the general constraints listed above, and any constraints listed in the [Limitations](#) section, in relation to other operations. Note that these additional constraints only apply when Reduction operations are used in the runtime fusion engine.

Table 16. Constraints for Reduction Operations

Attribute	Requirement
Tensor data type for CUDNN_ATTR_OPERATION_REDUCTION_YDESC	CUDNN_DATA_FLOAT
CUDNN_ATTR_REDUCTION_COMP_TYPE	CUDNN_DATA_FLOAT
Tensor layout for CUDNN_ATTR_OPERATION_REDUCTION_XDESC and CUDNN_ATTR_OPERATION_REDUCTION_YDESC	NHWC/NDHWC/BMN fully packed
CUDNN_ATTR_REDUCTION_OPERATOR	CUDNN_REDUCE_TENSOR_ADD, CUDNN_REDUCE_TENSOR_MIN, and CUDNN_REDUCE_TENSOR_MAX

Table 17. Supported Reduction Patterns

Reduction Operation	Reduction Pattern	
	Input	Output
Standalone reduction operation	[N, C, H, W]	[N, 1, H, W]
		[1, C, 1, 1]
		[1, 1, 1, 1]
Reduction fused after convolution backward filter gradient	[N, K, P, Q]	[N, 1, P, Q]
		[1, K, 1, 1]
		[1, 1, 1, 1]
Reduction fused after convolution backward data gradient	[N, C, H, W]	[N, 1, H, W]
		[1, C, 1, 1]
		[1, 1, 1, 1]
Reduction fused after convolution backward filter gradient	[K, C, R, S]	[K, 1, 1, 1]
		[1, C, R, S]
		[1, 1, 1, 1]
Reduction fused after matrix multiplication operation	[B, M, N]	[B, M, 1]
		[B, 1, N]

10.2.2.3.6. ResampleFwd

This operation represents resampling of the spatial dimensions of an image to a desired value. Resampling is supported in both directions, upsampling and downsampling. Downsampling represents the standard operation of pooling, commonly used in convolutional neural networks. Refer to the

[CUDNN_BACKEND_OPERATION_RESAMPLE_FWD_DESCRIPTOR](#) section for more information and general constraints.

The following are constraints for Resample operations, in addition to the general constraints listed above, and any constraints listed in the [Limitations](#) section, in relation to other operations. Note that these additional constraints only apply when Resample forward operations are used in the runtime fusion engine.

We allow a choice amongst four modes for resample. All modes have the following common support specifications:

- ▶ Supported layout: NHWC or NDHWC
- ▶ Spatial dimensions supported: 2 or 3
- ▶ Input dimensions supported: 4 or 5
- ▶ If specified, the index tensor dimension should be equal to the response tensor dimension.

There are some mode specific restrictions also. The following tables list the values that are allowed for particular parameters. For the parameters not listed, we allow any value which is mathematically correct.

The following downsampling modes are supported:

- ▶ CUDNN_RESAMPLE_AVGPOOL_INCLUDE_PADDING
- ▶ CUDNN_RESAMPLE_AVGPOOL_EXCLUDE_PADDING
- ▶ CUDNN_RESAMPLE_MAXPOOL

Table 18. Specific Restrictions for the Downsampling Modes

Attribute	Average Pooling	Max Pooling
CUDNN_ATTR_RESAMPLE_PADDING	CUDNN_ZERO_PAD	CUDNN_NEG_INF_PAD
CUDNN_ATTR_OPERATION_RESAMPLE_FWD_ALPHA	1.0	1.0
CUDNN_ATTR_OPERATION_RESAMPLE_FWD_BETA	0.0	0.0
CUDNN_ATTR_RESAMPLE_COMP_TYPE	CUDNN_DATA_FLOAT	CUDNN_DATA_FLOAT

For the upsampling modes, CUDNN_RESAMPLE_NEAREST is not supported for any combination of parameters. CUDNN_RESAMPLE_BILINEAR has the following support specifications.

Table 19. Specific Restrictions for Upsampling Mode

CUDNN_RESAMPLE_BILINEAR

Attribute	Bilinear
Input dimensions	Equal to 0.5 x output dimensions
CUDNN_ATTR_RESAMPLE_PRE_PADDINGS	0.5
CUDNN_ATTR_RESAMPLE_POST_PADDINGS	1

Attribute	Bilinear
CUDNN_ATTR_RESAMPLE_STRIDES	0.5
CUDNN_ATTR_RESAMPLE_WINDOW_DIMS	2
Data type for CUDNN_ATTR_OPERATION_RESAMPLE_FWD_XDESC and CUDNN_ATTR_OPERATION_RESAMPLE_FWD_YDESC	CUDNN_DATA_FLOAT
CUDNN_ATTR_RESAMPLE_COMP_TYPE	CUDNN_DATA_FLOAT
CUDNN_ATTR_OPERATION_RESAMPLE_FWD_ALPHA	1.0
CUDNN_ATTR_OPERATION_RESAMPLE_FWD_BETA	0.0
CUDNN_ATTR_RESAMPLE_PADDING_MODE	CUDNN_EDGE_VAL_PAD

10.2.2.3.6.1. Resampling Index Tensor Dump for Training


For max-pooling resampling mode, an index tensor can be provided to be used as a mask for backpropagation.

Values in the index tensors are:

- Zero-indexed row-major position of maximum value of input tensor in the resampling window.
- In case of multiple input pixels with maximum value, the first index in a left-to-right top-to-bottom scan is selected.

Example of index element selection:

Figure 24. Values In the Index Tensors



Select an appropriate element size for the index tensor. As a reference, any element size such that the maximum zero-indexed window position fits should be sufficient.

10.2.2.3.7. ResampleBwd

This operation represents backward resampling of the spatial dimensions of an output response to a desired value. Resampling is supported in both directions, upsampling and downsampling. Backwards downsampling represents the standard operation of backward pooling, commonly used in convolutional neural networks. Refer to the [CUDNN_BACKEND_OPERATION_RESAMPLE_BWD_DESCRIPTOR](#) section for more information and general constraints.

The following are constraints for Resample backward operations, in addition to the general constraints listed above, and any constraints listed in the [Limitations](#) section, in relation to other operations. Note that these additional constraints only apply when Resample backward operations are used in the runtime fusion engine.

We allow a choice amongst four modes for resample. All modes have the following common support specifications:

- ▶ Supported layout: NHWC or NDHWC
- ▶ Spatial dimensions supported: 2 or 3
- ▶ Input dimensions supported: 4 or 5
- ▶ The index tensor dimensions should be equal to the input gradient tensor dimensions.

Index tensor should be provided for only max pooling mode, and should adhere to the format described in the [Resampling Index Tensor Dump for Training](#) section.

There are some mode specific restrictions also. The following tables list the values that are allowed for particular parameters. For the parameters not listed, we allow any value which is mathematically correct.

The following backward downsampling modes are supported:

- ▶ CUDNN_RESAMPLE_AVGPOOL_INCLUDE_PADDING
- ▶ CUDNN_RESAMPLE_AVGPOOL_EXCLUDE_PADDING
- ▶ CUDNN_RESAMPLE_MAXPOOL

Table 20. Specific Restrictions for the Backwards Downsampling Modes

Attribute	Average Pooling	Max Pooling
CUDNN_ATTR_RESAMPLE_PADDING	CUDNN_ZERO_PAD	CUDNN_NEG_INF_PAD
CUDNN_ATTR_OPERATION_RESAMPLE_BWD_ALPHA		1.0
CUDNN_ATTR_OPERATION_RESAMPLE_BWD_BETA		0.0
CUDNN_ATTR_RESAMPLE_COMP_TYPE	CUDNN_DATA_FLOAT	CUDNN_DATA_FLOAT

Backward upsampling modes are currently not supported.

10.2.3. Pre-compiled Specialized Engines

The pre-compiled specialized engines target and optimize for a specialized graph pattern with a ragged support surface. Because of this targeting, these graphs do not require runtime compilation.

In most cases, the specialized patterns are just special cases of the generic patterns used in the runtime fusion engine, but there are some cases where the specialized pattern does not fit any of the generic patterns. If your graph pattern matches a specialized pattern, you will get at least a pattern matching engine, and you might also get a runtime fusion engine as another option.


Currently, the following patterns are supported by the pattern matching engines. Some nodes are optional. Optional nodes are indicated by dashed outlines.

10.2.3.1. ConvBNfprop

In [Figure 25](#), the `ConvBNfprop` pattern is illustrated. Its restrictions and options include:

1. The three pointwise nodes scale, bias, and ReLU are optional.
2. X , Z , W , s_1 , b_1 must all be of FP16 data type.
3. Z needs to be of shape $[N, C, H, W]$ with NHWC packed layout.
4. W needs to be of shape $[K, C, R, S]$ with KRSC packed layout.
5. s_1 , b_1 need to be of shape $[1, C, 1, 1]$ with NHWC packed layout.
6. Only ReLU activation is supported.
7. All of the intermediate tensors need to be virtual, except, Y needs to be non-virtual.
8. I/O pointers should be 16 bytes aligned.

Figure 25. The pre-compiled `ConvBNfprop` engine fuses several pointwise operations with `ConvolutionFwd` and `GenStats`.




10.2.3.2. ConvBNwgrad

In [Figure 26](#), the `ConvBNwgrad` pattern is illustrated. Its restrictions and options include:

1. The three pointwise operations are all optional, as indicated by the dashed outlines.
2. Only ReLU activation is supported.
3. X , s_1 , b_1 , and dy must all be of FP16 datatype.
4. I/O pointers should be 16 bytes aligned.

Figure 26. The `ConvBNwgrad` pre-compiled engine fuses several (optional) pointwise operations with `ConvolutionBwFilter`.




10.2.3.3. ConvBiasAct

In the following figure, the `ConvBiasAct` pattern is illustrated. Its restrictions and options include:

1. α_1 and α_2 need to be scalars.
2. The activation node is optional.
3. The size of the bias tensor should be $[1, K, 1, 1]$.
4. Internal conversions are not supported. That is, the virtual output between nodes need to have the same data type as the node's compute type, which should be the same as the epilog type of the convolution node.
5. There are some restrictions on the supported combination of data types, which can be found in the API Reference (refer to [cudnnConvolutionBiasActivationForward\(\)](#)).

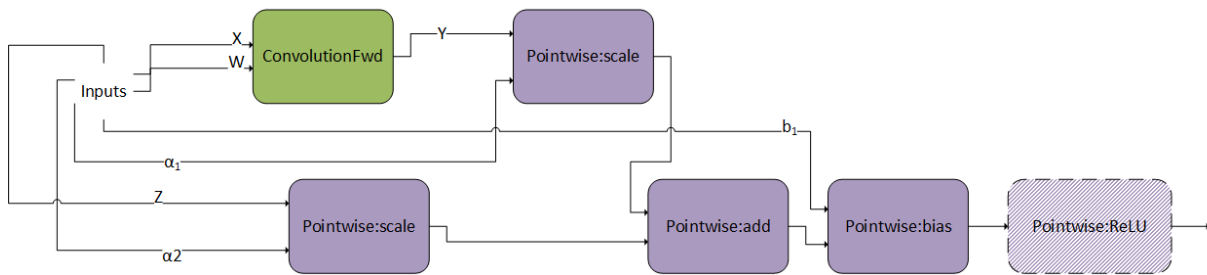
Figure 27. `ConvBiasAct`, another pre-compiled engine, fuses `ConvolutionFwd` with several pointwise operations.



10.2.3.4. ConvScaleBiasAct

In the following figure, the `ConvScaleBiasAct` pattern is illustrated. Its restrictions and options include:

1. α_1 and α_2 and b_2 should have the same data type/layout and can only be FP32.
2. X, W, and Z can only be int8x4 or int8x32.
3. The size of the bias tensor should be $[1, K, 1, 1]$.
4. Internal conversions are not supported. Meaning, "virtual output" between nodes needs to be the same as their compute type.
5. Currently, `Pointwise:ReLU` is the only optional pointwise node.

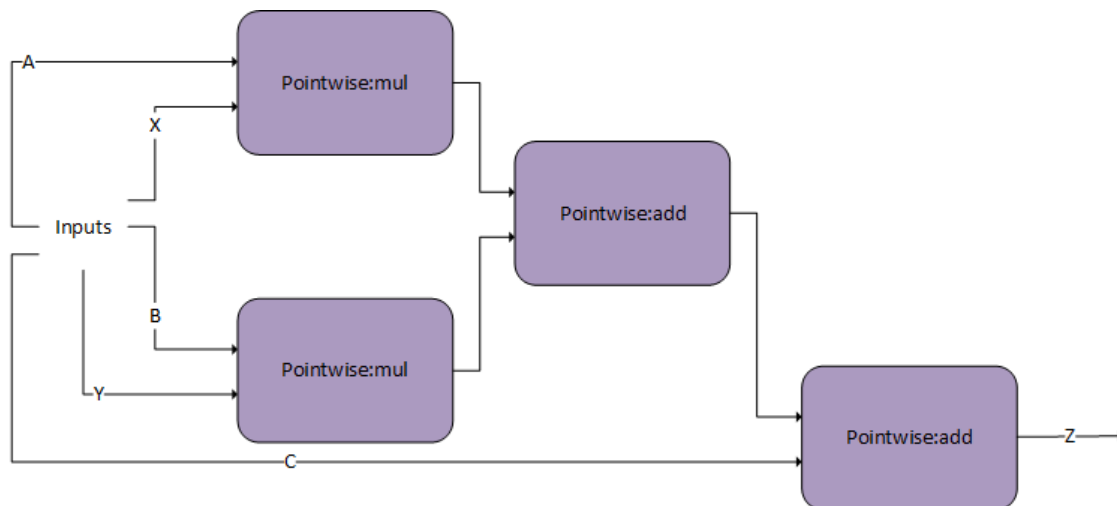
Figure 28. The pre-compiled engine, `ConvScaleBiasAct`

This pattern is very similar as `ConvBiasAct`. The difference is that here, the scales α_1 and α_2 are tensors, not scalars. If they are scalars, this pattern becomes a normal `ConvBiasAct`.

10.2.3.5. `dBNAppl`

In `dBNAppl`, the `dBNAppl` pattern is illustrated. Its restrictions and options include:

1. One of the inputs to the `mul` nodes and the input to the final `add` node must be of FP32 datatype (A, B, C).
2. The other inputs to the `mul` nodes (X and Y) must be of FP16 data type.
3. X, Y and Z are 4D tensors – [N,C,H,W] with NHWC packed layout.
4. A, B, C are 1D tensors – [1,C,1,1] with NHWC packed layout.
5. Channel C should be a multiple of 16 for all the tensors.
6. Tensors A and B should be attached to the B port of the `mul` nodes; tensors X and Y should be attached to the X port.

Figure 29. The pre-compiled engine, `dBNAppl`

The pattern implements a simple linear combination:

► $Z = A * X + B * Y + C$

10.2.3.6. DualdBNApply

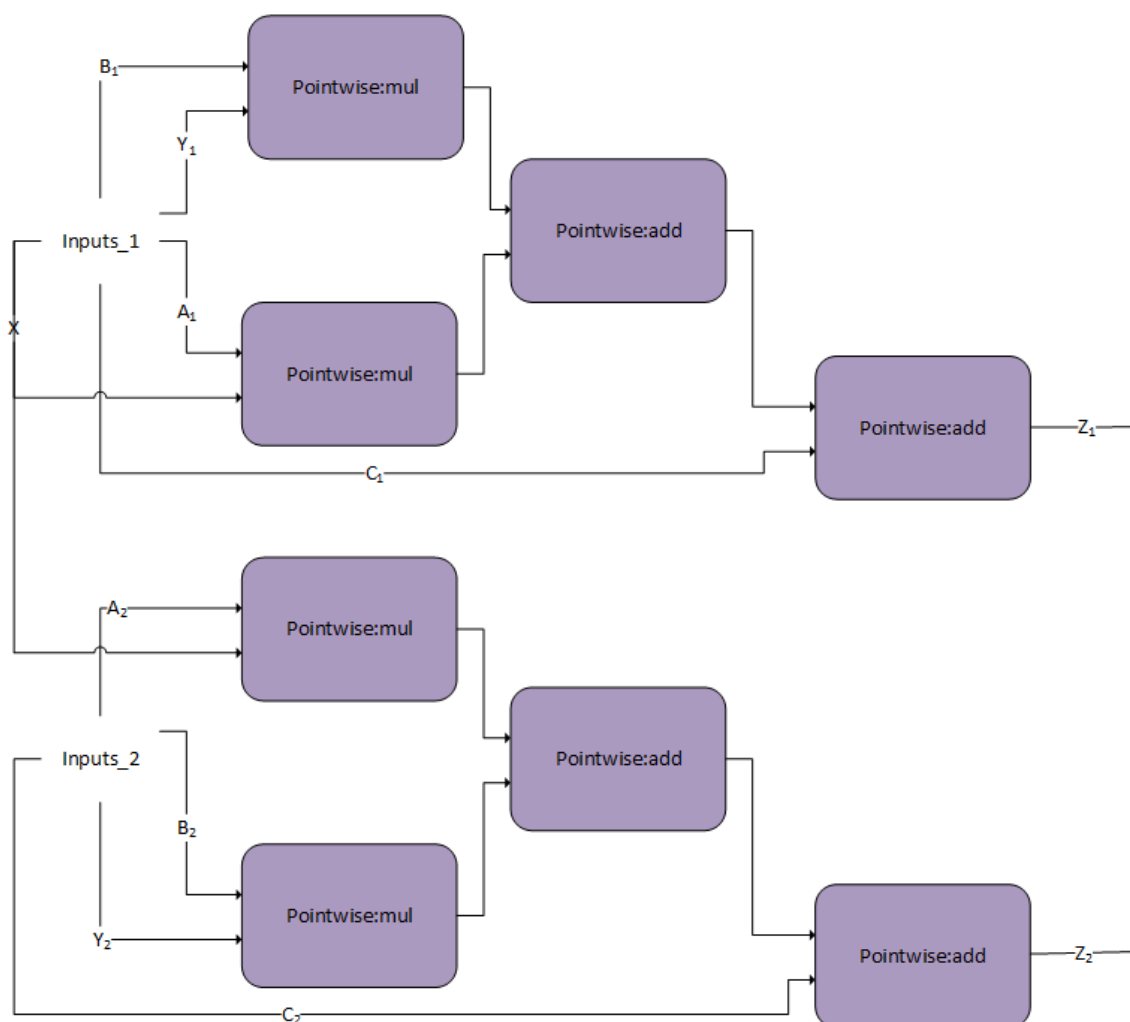
In [Figure 25](#), the `DualdBNApply` pattern is illustrated. Its restrictions and options include:

1. One tensor X is shared between the two linear combinations.
2. Five tensors, X, Y_1, Y_2, Z_1, Z_2 are 4D tensors $[N, C, H, W]$ with NHWC packed layout.
3. Six tensors $A_1, A_2, B_1, B_2, C_1, C_2$ are 1D tensors $[1, C, 1, 1]$.
4. Channel C should be a multiple of 16 for all the tensors.

In essence, `DualdBNApply` runs the previous pattern, `dBNApply` twice, as two subgraphs. However, both subgraphs share one input tensor, X .

Note that for visibility purposes, the Inputs block is split into `Inputs_1` and `Inputs_2`. This has no semantic meaning.

Figure 30. The `DualdBNApply` engine



This pattern implements two linear combinations:

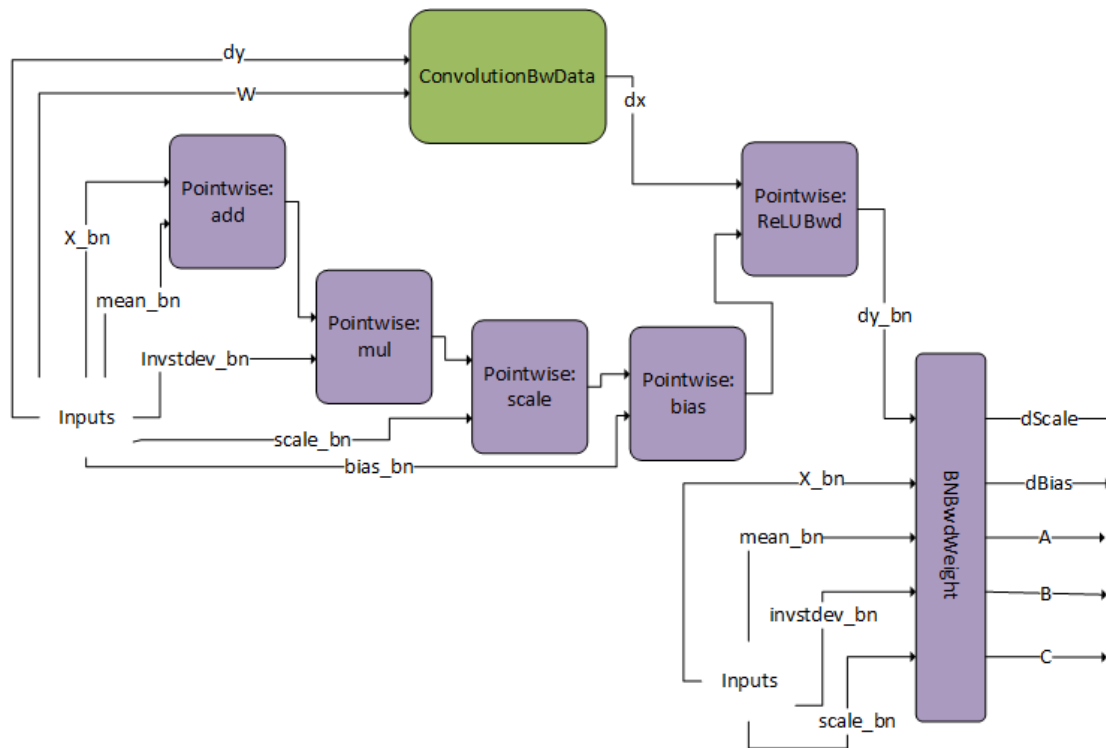
- $Z_1 = A_1 * X + B_1 * Y_1 + C_1$
- $Z_2 = A_2 * X + B_2 * Y_2 + C_2$

10.2.3.7. DgradDreluBNBwdWeight

In [Figure 31](#), the `DgradDreluBNBwdWeight` pattern is illustrated. Its restrictions and options include:

1. Dgrad input `dy` and `W` are of FP16 datatypes.
2. Batch norm fwd inputs, `X_bn` is of FP16 datatype while the other tensors `mean_bn`, `invstd_dev_bn`, `scale_bn`, and `bias_bn` are FP32.
3. Outputs: `dScale`, `dBias`, `A`, `B`, `C` are of FP32 data type.
4. All pointers are 16 byte aligned.
5. Only supported on NVIDIA Ampere Architecture GPUs.

Figure 31. `DgradDreluBNBwdWeight` is a pre-compiled engine that can be used in conjunction with the `dBNAppl` pattern to compute the backwards path of batch norm.



The `BNBwdWeight` operation takes in five inputs: `X_bn`, `mean_bn`, `invstddev_bn`, `scale_bn`, and `dy_bn` (that is, the output from the `ReLUBwd` node).

It produces five outputs: gradients of the batch norm scale and bias params, `dScale`, `dBias`, and coefficients `A`, `B`, `C`. Note that for illustration purposes, the inputs are duplicated. The inputs on the left and right are however exactly the same.

This pattern is typically used in the computation of the Batch Norm Backward Pass.

When computing the backward pass of batch norm, `dScale`, `dBias`, and `dx_bn` are needed. The `DgradDreluBnBwdWeight` pattern computes the former two. Using the generated A, B, and C we can use the `dBNApplly` pattern above to compute `dX`, the input gradient, as follows $dx_bn = A * dy_bn + B * X_bn + C$.

Note that this pattern is used in combination with the forward pass, the `ConvBNfprop` pattern. Because of performance reasons, the output of Batch Norm `Y_bn`, which was calculated in `ConvBNfprop` (output of scale-bias), needs to be recalculated by `DgradDreluBnBwdWeight`. The pointwise add node subtracts `mean_bn` from `X_bn`, hence the `alpha2` parameter for that node should be set to `-1`.

10.2.4. Mapping with Backend Descriptors

For readability, the operations used in this section are abbreviated. The mapping with the actual backend descriptors can be found in this table:

Table 21. Notations and Backend Descriptors

Notation used in this section	Backend descriptor
<code>Pointwise:scale</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_MUL</code> and with operand B broadcasting into operand X
<code>Pointwise:bias</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_ADD</code> and with operand B broadcasting into operand X
<code>Pointwise:add</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_ADD</code> and with operand B with same dimensions as X
<code>Pointwise:mul</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_MUL</code> and with operand B with same dimensions as X
<code>Pointwise:ReLU</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_RELU_FWD</code>
<code>Pointwise:ReLUbwd</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_RELU_BWD</code>
<code>Pointwise:tanh</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_TANH_FWD</code>
<code>Pointwise:sigmoid</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_SIGMOID_FWD</code>
<code>Pointwise:ELU</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with mode <code>CUDNN_POINTWISE_ELU_FWD</code>
<code>Pointwise:{ReLU,tanh,sigmoid,ELU}</code>	<code>CUDNN_BACKEND_OPERATION_POINTWISE_DESCRIPTOR</code> with one of the following modes: <code>CUDNN_POINTWISE_RELU_FWD,</code> <code>CUDNN_POINTWISE_TANH_FWD,</code> <code>CUDNN_POINTWISE_SIGMOID_FWD,</code> <code>CUDNN_POINTWISE_ELU_FWD</code>

Notation used in this section	Backend descriptor
MatMul	CUDNN_BACKEND_OPERATION_MATMUL_DESCRIPTOR
ConvolutionFwd	CUDNN_BACKEND_OPERATION_CONVOLUTION_FORWARD_DESCRIPTOR
ConvolutionBwFilter	CUDNN_BACKEND_OPERATION_CONVOLUTION_BACKWARD_FILTER_DESCRIPTOR
ConvolutionBwData	CUDNN_BACKEND_OPERATION_CONVOLUTION_BACKWARD_DATA_DESCRIPTOR
GenStats	CUDNN_BACKEND_OPERATION_GEN_STATS_DESCRIPTOR
ResampleFwd	CUDNN_BACKEND_OPERATION_RESAMPLE_FWD_DESCRIPTOR
GenStats	CUDNN_BACKEND_OPERATION_GEN_STATS_DESCRIPTOR
Reduction	CUDNN_BACKEND_OPERATION_REDUCTION_DESCRIPTOR
BnBwdWeight	CUDNN_BACKEND_OPERATION_BN_BWD_WEIGHTS_DESCRIPTOR
BOOLEAN/packed-BOOLEAN	<p>CUDNN_DATA_BOOLEAN: As described in the API Reference, this type implies that eight boolean values are packed in a single byte, with the lowest index on the right (that is, least significant bit).</p> <p>packed-BOOLEAN and BOOLEAN are used interchangeably, where the former is used to emphasize and remind the user about the semantics.</p>
INT8	CUDNN_DATA_INT8
FP8	CUDNN_DATA_FP8_E4M3 or CUDNN_DATA_FP8_E5M2
FP16	CUDNN_DATA_HALF
BF16	CUDNN_DATA_BFLOAT16
FP32	CUDNN_DATA_FLOAT
TF32	A tensor core operation mode used to accelerate floating point convolutions or matmuls. This can be used for an operation with compute type CUDNN_DATA_FLOAT, on Ampere or later and be disabled with NVIDIA_TF32_OVERRIDE=1.

Chapter 11. Troubleshooting

The following sections help answer the most commonly asked questions regarding typical use cases.

11.1. Error Reporting And API Logging

The cuDNN error reporting and API logging is a utility for recording the cuDNN API execution and error information. For each cuDNN API function call, all input parameters are reported in the API logging. If errors occur during the execution of the cuDNN API, a traceback of the error conditions can also be reported to help troubleshooting. This functionality is disabled by default, and can be enabled using the methods described in the later part of this section through three logging severity levels: `CUDNN_LOGINFO_DBG`, `CUDNN_LOGWARN_DBG` and `CUDNN_LOGERR_DBG`.

The log output contains variable names, data types, parameter values, device pointers, process ID, thread ID, cuDNN handle, CUDA stream ID, and metadata such as time of the function call in microseconds.

For example, when the severity level `CUDNN_LOGINFO_DBG` is enabled, the user will receive the API loggings, such as:

```
cuDNN (v8300) function cudnnSetActivationDescriptor() called:
  mode: type=cudnnActivationMode_t; val=CUDNN_ACTIVATION_RELU (1);
  reluNanOpt: type=cudnnNanPropagation_t; val=CUDNN_NOT_PROPAGATE_NAN (0);
  coef: type=double; val=1000.000000;
Time: 2017-11-21T14:14:21.366171 (0d+0h+1m+5s since start)
Process: 21264, Thread: 21264, cudnn_handle: NULL, cudnn_stream: NULL.
```

Starting in cuDNN 8.3.0, when the severity level `CUDNN_LOGWARN_DBG` or `CUDNN_LOGERR_DBG` are enabled, the log output additionally reports an error traceback such as the example below (currently only cuDNN version 8 graph APIs and legacy convolution APIs are using this error reporting feature). This traceback reports the relevant error/warning conditions, aiming to provide the user hints for troubleshooting purposes. Within the traceback, each message may have their own severity and will only be reported when the respective severity level is enabled. The traceback messages are printed in the reverse order of the execution so the messages at the top will be the root cause and tend to be more helpful for debugging.

```
cuDNN (v8300) function cudnnBackendFinalize() called:
  Info: Traceback contains 5 message(s)
    Error: CUDNN_STATUS_BAD_PARAM; reason: out <= 0
    Error: CUDNN_STATUS_BAD_PARAM; reason: is_valid_spatial_dim(xSpatialDimA[dim],
wSpatialDimA[dim], ySpatialDimA[dim], cDesc.getPadLowerA()[dim], cDesc.getPadUpperA()[dim],
cDesc.getStrideA()[dim], cDesc.getDilationA()[dim])
```

```
Error: CUDNN_STATUS_BAD_PARAM; reason: is_valid_convolution(xDesc, wDesc, cDesc,
yDesc)
Error: CUDNN_STATUS_BAD_PARAM; reason: convolution_init(xDesc, wDesc, cDesc, yDesc)
Error: CUDNN_STATUS_BAD_PARAM; reason: finalize_internal()
Time: 2021-10-05T17:11:07.935640 (0d+0h+0m+15s since start)
Process=87720; Thread=87720; GPU=NULL; Handle=NULL; StreamId=NULL.
```

There are two methods, as described below, to enable the error/warning reporting and API logging. For convenience, the log output can be handled by the built-in default callback function, which will direct the output to a log file or the standard I/O as designated by the user. The user may also write their own callback function to handle this information programmatically, and use the [cudnnSetCallback\(\)](#) to pass in the function pointer of their own callback function.

Method 1: Using Environment Variables

To enable API logging using environment variables, follow these steps:

- ▶ Decide which logging severity levels to include from these three options: CUDNN_LOGINFO_DBG, CUDNN_LOGWARN_DBG, CUDNN_LOGERR_DBG. The logging severity levels are independent of each other. Any combination of them is valid.
- ▶ Set the environment variables CUDNN_LOGINFO_DBG or CUDNN_LOGWARN_DBG or CUDNN_LOGERR_DBG to 1, and
- ▶ Set the environment variable CUDNN_LOGDEST_DBG to one of the following:
 - ▶ stdout, stderr, or a user-desired file path, for example, /home/userName1/log.txt.
- ▶ Include the conversion specifiers in the file name. For example:
 - ▶ To include date and time in the file name, use the date and time conversion specifiers: log_%Y_%m_%d_%H_%M_%S.txt. The conversion specifiers will be automatically replaced with the date and time when the program is initiated, resulting in log_2017_11_21_09_41_00.txt.
 - ▶ To include the process id in the file name, use the %i conversion specifier: log_%Y_%m_%d_%H_%M_%S_%i.txt for the result: log_2017_11_21_09_41_00_21264.txt when the process id is 21264. When you have several processes running, using the process id conversion specifier will prevent these processes from writing to the same file at the same time.



Note: The supported conversion specifiers are similar to the `strftime` function.

If the file already exists, the log will overwrite the existing file.



Note: These environmental variables are only checked once at the initialization. Any subsequent changes in these environmental variables will not be effective in the current run. Also note that these environment settings can be overridden by Method 2 below.

Refer to [Table 22](#) for the impact on the performance of API logging using environment variables. The `CUDNN_LOG{INFO,WARN,ERR}_DBG` notation in the table header means the conclusion is applicable to either one of the environment variables.

Table 22. API Logging Using Environment Variables

Environment variables	<code>CUDNN_LOG{INFO,WARN,ERR}_DBG</code>	<code>CUDNN_LOG{INFO,WARN,ERR}_DBG=1</code>
<code>CUDNN_LOGDEST_DBG</code> not set	No logging output No performance loss	No logging output No performance loss
<code>CUDNN_LOGDEST_DBG=NULL</code>	No logging output No performance loss	No logging output No performance loss
<code>CUDNN_LOGDEST_DBG=stdout</code> or <code>stderr</code>	No logging output No performance loss	Logging to <code>stdout</code> or <code>stderr</code> Some performance loss
<code>CUDNN_LOGDEST_DBG=filename.txt</code>	No logging output No performance loss	Logging to <code>filename.txt</code> Some performance loss

Method 2: Using the API

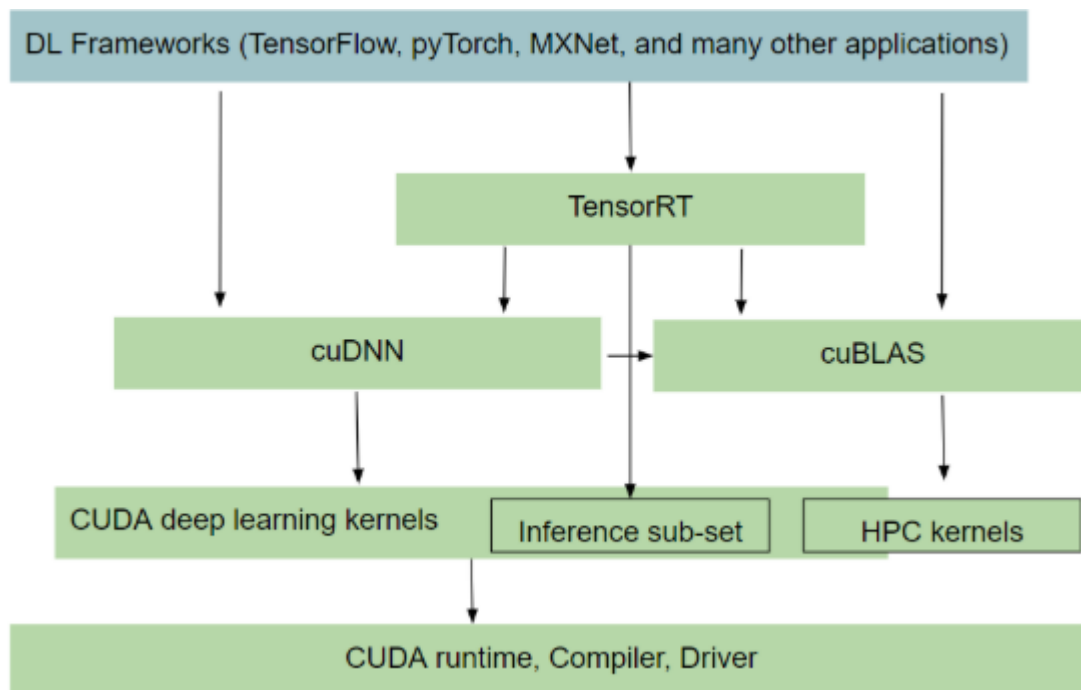
Method 2: To use API function calls to enable API logging, refer to the API description of [`cudaSetCallback\(\)`](#) and [`cudaGetCallback\(\)`](#).

11.2. FAQs

Q: Where in the software stack does cuDNN sit? What is the interaction between CUDA, cuDNN, and TensorRT?

A: The following graphic shows how cuDNN relates to other software in the stack.

Figure 32. Software stack with cuDNN.



Q: I'm not sure if I should use cuDNN for inference or training. How does it compare with TensorRT?

A: cuDNN provides the building blocks for common routines such as convolution, pooling, activation and RNN/LSTMs. You can use cuDNN for both training and inference. However, where it differs from TensorRT is that the latter (TensorRT) is a programmable inference accelerator; just like a framework. TensorRT sees the whole graph and optimizes the network by fusing/combining layers and optimizing kernel selection for improved latency, throughput, power efficiency and for reducing memory requirements.

A rule of thumb you can apply is to check out TensorRT, see if it meets your inference needs, if it doesn't, then look at cuDNN for a closer, more in-depth perspective.

Q: How does heuristics in cuDNN work? How does it know what is the optimal solution for a given problem?

A: NVIDIA actively monitors the Deep Learning space for important problem specifications such as commonly used models. The heuristics are produced by sampling a portion of these problem specifications with available computational choices. Over time, more models are discovered and incorporated into the heuristics.

Q: Is cuDNN going to support running arbitrary graphs?

A: No, we don't plan to become a framework and execute the whole graph one op at a time. At this time, we are focused on a subgraph given by the user, where we try to

produce an optimized fusion kernel. We will document the rules regarding what can be fused and what cannot. The goal is to support general and flexible fusion, however, it will take time and there will be limits in what it can do in the cuDNN version 8.0.0 launch.

Q: What's the difference between TensorRT, TensorFlow/XLA's fusion, and cuDNN's fusion?

A: TensorRT and TensorFlow are frameworks; they see the whole graph and can do global optimization, however, they generally only fuse pointwise ops together or pattern match to a limited set of pre-compiled fixed fusion patterns like conv-bias-relu. On the other hand, cuDNN targets a subgraph, but can fuse convolutions with pointwise ops, thus providing potentially better performance. CuDNN fusion kernels can be utilized by TensorRT and TensorFlow/XLA as part of their global graph optimization.

Q: Can I write an application calling cuDNN directly?

A: Yes, you can call the C/C++ API directly. Usually, data scientists would wait for framework integration and use the Python API which is more convenient. However, if your use case requires better performance, you can target the cuDNN API directly.

Q: How does mixed precision training work?

A: Several components need to work together to make mixed precision training possible. CuDNN needs to support the layers with the required datatype config and have optimized kernels that run very fast. In addition, there is a module called automatic mixed precision (AMP) in frameworks which intelligently decides which op can run in a lower precision without affecting convergence and minimize the number of type conversions/transposes in the entire graph. These work together to give you speed up. For more information, see [Mixed Precision Numerical Accuracy](#).

Q: How can I pick the fastest convolution kernels with cuDNN version 8.0.0?

A: In the API introduced in cuDNN v8, convolution kernels are grouped by similar computation and numerical properties into engines. Every engine has a queryable set of performance tuning knobs. A computation case such as a convolution operation graph can be computed using different valid combinations of engines and their knobs, known as an engine configuration. Users can query an array of engine configurations for any given computation case ordered by performance, from fastest to slowest according to cuDNN's own heuristics. Alternately, users can generate all possible engine configurations by querying the engine count and available knobs for each engine. This generated list could be used for auto-tuning or the user could create their own heuristics.

Q: Why is cuDNN version 8.0 convolution API call much slower on the first call than subsequent calls?

A: Due to the library split, cuDNN version 8.0 API will only load the necessary kernels on the first API call that requires it. In previous versions, this load would have been observed in the first cuDNN API call that triggers CUDA context initialization, typically `cudaCreate()`. In version 8.0, this is delayed until the first sub-library call that triggers CUDA context initialization. Users who desire to have CUDA context preloaded can call the new `cudaCnnInferVersionCheck()` API (or its related cousins), which has the side effect of initializing a CUDA context. This will reduce the run time for all subsequent API calls.

Q: How do I build the cuDNN version 8.0.0 split library?

A: cuDNN v8.0 library is split into multiple sub-libraries. Each library contains a subset of the API. Users can link directly against the individual libraries or link with a `dlopen` layer which follows a plugin architecture.

To link against an individual library, users can directly specify it and its dependencies on the linker command line. For example, for infer libraries: `-lcudnn_adv_infer, -lcudnn_cnn_infer`, or `-lcudnn_ops_infer`.

For all libraries, `-lcudnn_adv_train, -lcudnn_cnn_train, -lcudnn_ops_train, -lcudnn_adv_infer, -lcudnn_cnn_infer`, and `-lcudnn_ops_infer`.

The dependency order is documented in the cuDNN [8.0.0 Preview Release Notes](#) and the [NVIDIA cuDNN API Reference](#).

Alternatively, the user can continue to link against a shim layer (`-libcudnn`) which can `dlopen` the correct library that provides the implementation of the function. When the function is called for the first time, the dynamic loading of the library takes place.

Linker argument:

```
-libcudnn
```

Q: What are the new APIs in cuDNN version 8.0.0?

A: The new cuDNN APIs are listed in the cuDNN 8.0.0 Release Notes as well as in the [API Changes For cuDNN 8.0.0](#).

11.3. Support

Support, resources, and information about cuDNN can be found online at <https://developer.nvidia.com/cudnn>. This includes downloads, webinars, [NVIDIA Developer Forums](#), and more.

We appreciate all types of feedback. Consider posting on the forums with questions, comments, and suspected bugs that are appropriate to discuss publicly. cuDNN-related posts are reviewed by the cuDNN engineering team, and internally we will file bugs where appropriate. It's helpful if you can paste or attach an [Error Reporting And API Logging](#) to help us reproduce.

External users can also file bugs directly by following these steps:

1. Register for the [NVIDIA Developer website](#).
2. Log in to the developer site.
3. Click on your name in the upper right corner.
4. Click My account > My Bugs and select Submit a New Bug.
5. Fill out the bug reporting page. Be descriptive and if possible, provide the steps that you are following to help reproduce the problem. If possible, paste or attach an [API log](#).
6. Click Submit a bug.

Chapter 12. Acknowledgments

Some of the cuDNN library routines were derived from code developed by others and are subject to the following:

12.1. University of Tennessee

```
Copyright (c) 2010 The University of Tennessee.
```

```
All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without  
modification, are permitted provided that the following conditions are  
met:
```

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer listed in this license in the documentation and/or other materials provided with the distribution.
- * Neither the name of the copyright holders nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS  
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT  
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR  
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT  
OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,  
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT  
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,  
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY  
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT  
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE  
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

12.2. University of California, Berkeley

```
COPYRIGHT
```

```
All contributions by the University of California:  
Copyright (c) 2014, The Regents of the University of California (Regents)  
All rights reserved.
```

```
All other contributions:  
Copyright (c) 2014, the respective contributors  
All rights reserved.
```

Caffe uses a shared copyright model: each contributor holds copyright over their contributions to Caffe. The project versioning records all such contribution and copyright details. If a contributor wants to further mark their specific copyright on a particular contribution, they should indicate their copyright solely in the commit message of the change when it is committed.

LICENSE

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

CONTRIBUTION AGREEMENT

By contributing to the BVLC/caffe repository through pull-request, comment, or otherwise, the contributor releases their content to the license and copyright terms herein.

12.3. Facebook AI Research, New York

Copyright (c) 2014, Facebook, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name Facebook nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS

SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Additional Grant of Patent Rights

"Software" means fbcunn software distributed by Facebook, Inc.

Facebook hereby grants you a perpetual, worldwide, royalty-free, non-exclusive, irrevocable (subject to the termination provision below) license under any rights in any patent claims owned by Facebook, to make, have made, use, sell, offer to sell, import, and otherwise transfer the Software. For avoidance of doubt, no license is granted under Facebook's rights in any patent claims that are infringed by (i) modifications to the Software made by you or a third party, or (ii) the Software in combination with any software or other technology provided by you or a third party.

The license granted hereunder will terminate, automatically and without notice, for anyone that makes any claim (including by filing any lawsuit, assertion or other action) alleging (a) direct, indirect, or contributory infringement or inducement to infringe any patent: (i) by Facebook or any of its subsidiaries or affiliates, whether or not such claim is related to the Software, (ii) by any party if such claim arises in whole or in part from any software, product or service of Facebook or any of its subsidiaries or affiliates, whether or not such claim is related to the Software, or (iii) by any party relating to the Software; or (b) that any right in any patent claim of Facebook is invalid or unenforceable.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Arm

Arm, AMBA and Arm Powered are registered trademarks of Arm Limited. Cortex, MPCore and Mali are trademarks of Arm Limited. "Arm" is used to represent Arm Holdings plc; its operating company Arm Limited; and the regional subsidiaries Arm Inc.; Arm KK; Arm Korea Limited.; Arm Taiwan Limited; Arm France SAS; Arm Consulting (Shanghai) Co. Ltd.; Arm Germany GmbH; Arm Embedded Technologies Pvt. Ltd.; Arm Norway, AS and Arm Sweden AB.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

BlackBerry/QNX

Copyright © 2020 BlackBerry Limited. All rights reserved.

Trademarks, including but not limited to BLACKBERRY, EMBLEM Design, QNX, AVIAGE, MOMENTICS, NEUTRINO and QNX CAR are the trademarks or registered trademarks of BlackBerry Limited, used under license, and the exclusive rights to such trademarks are expressly reserved.

Google

Android, Android TV, Google Play and the Google Play logo are trademarks of Google, Inc.

Trademarks

NVIDIA, the NVIDIA logo, and BlueField, CUDA, DALI, DRIVE, Hopper, JetPack, Jetson AGX Xavier, Jetson Nano, Maxwell, NGC, Nsight, Orin, Pascal, Quadro, Tegra, TensorRT, Triton, Turing and Volta are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2017-2024 NVIDIA Corporation & affiliates. All rights reserved.

